

# **VIRTEX-5QV STATIC SEU CHARACTERIZATION SUMMARY**

Gary Swift  
Xilinx, Inc.  
San Jose, California

Gregory Allen  
Jet Propulsion Laboratory  
Pasadena, California

with help from other members of the  
Xilinx Radiation Test Consortium

## Document Revision History

Draft Release: July 26, 2011

Release Candidate: June 25, 2012

Official Release: July 15, 2012

In-progress revisions: Feb. 27, 2012

Second Release: May 17, 2013

Full editing and re-writing pass on all sections for clarity and correctness plus significant new material added, including proton test results.

A portion of the work reported here was carried out at the Jet Propulsion Laboratory, California Institute of Technology, and was sponsored by the National Aeronautics and Space Administration Electronic Parts and Packaging (NEPP) Program.

Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not constitute or imply its endorsement by the United States Government or the Jet Propulsion Laboratory, California Institute of Technology.

Copyright 2011-2013. All rights reserved.

# TABLE OF CONTENTS

	TABLE OF FIGURES .....	iv
	TABLE OF TABLES .....	v
1	VIRTEX-5 OVERVIEW .....	1
2	LATCHUP TESTING .....	3
3	STATIC TESTING .....	9
3.1	Experiment Setup.....	9
3.2	SEFI Results.....	11
3.2.1	<i>Experiment Considerations</i> .....	11
3.2.2	<i>SEFI Nomenclature and Taxonomy</i> .....	12
3.2.3	<i>SEFI Measurements and Data Fits</i> .....	14
3.3	Other Static Results.....	17
3.3.1	<i>Block Memory</i> .....	18
3.3.2	<i>User Flip-Flops</i> .....	22
3.3.3	<i>DSP Registers</i> .....	24
3.4	Orbital Rate Calculations.....	27
4	CONFIGURATION (DUAL-NODE) RESULTS .....	28
4.1	Test Chip Static Results .....	28
4.2	Configuration Cell Direct-Upset Static Results.....	33
4.3	SET-triggered FPGA Results.....	34
4.4	Orbital Rate Calculations .....	36
5	CONCLUSION .....	37
6	REFERENCES .....	37
7	APPENDIX.....	38

## TABLE OF FIGURES

<b>Figure 1a.</b> Current and temperature for s/n: 592 during latchup testing: angle=50 degrees, fluence = $2.0 \times 10^7$ gold ions per sq.cm, effective LET=135.7 MeV per mg/cm <sup>2</sup> .....	5
<b>Figure 1b.</b> Current and temperature for s/n:515 during latchup testing: angle=50 degrees, fluence = $2.0 \times 10^7$ gold ions per sq.cm, effective LET=145.2 MeV per mg/cm <sup>2</sup> .....	6
<b>Figure 1c.</b> Current and temperature for s/n:514 during latchup testing: angle=50 degrees, fluence = $2.0 \times 10^7$ gold ions per sq.cm, effective LET=145.5 MeV per mg/cm <sup>2</sup> .....	7
<b>Figure 2.</b> Current and temperature for s/n: 592 during latchup testing: angle=50 degrees, fluence < $2.8 \times 10^5$ gold ions per sq.cm, effective LET=135.7 MeV per mg/cm <sup>2</sup> .....	8
<b>Figure 3.</b> SEU Test Hardware Setup .....	10
<b>Figure 4a.</b> Virtex-5QV Susceptibility to ALL SEFIs from Heavy Ion Strikes .....	15
<b>Figure 4b.</b> Virtex-5QV SEFI Susceptibility due to Protons .....	16
<b>Figure 5.</b> Virtex-5QV BRAM Susceptibility due to Heavy Ion Strikes.....	18
<b>Figure 6.</b> Virtex-5QV BRAM Susceptibility due to Proton Strikes.....	18
<b>Figure 7.</b> Virtex-5QV BRAM Susceptibility Angle Experiment v. HICUP Model Fit .....	19
<b>Figure 8.</b> Upset cross section vs. Effective LET for He Irradiation .....	21
<b>Figure 9.</b> Low Energy Proton LET Spectrum .....	21
<b>Figure 10a.</b> Virtex-5QV User Flip-Flop (with filters OFF) Pseudo-static (1.5 MHz) Susceptibility to Heavy Ions .....	22
<b>Figure 10b.</b> Measured Virtex-5QV User Flip-Flop Pseudo-static (1.5 MHz) Susceptibility to Protons and Weibull Fits. The data is sparse, to say the least, esp. with filters ON so a 64 MeV, 200 MHz “anchor” point is shown.....	23
<b>Figure 11a.</b> Virtex-5QV DSP A/B-, C-, and P-Register Susceptibility to Heavy Ions; C-Register data shown. Note that each DSP has one M register and three others: A/B, C, and P although register use is optional; therefore, duty cycles vary by design and op code. ....	24
<b>Figure 11b.</b> Virtex-5QV DSP M-Register Susceptibility to Heavy Ions. Note that each DSP has one M register and three others: A/B, C, and P although register use is optional; therefore, the duty cycle varies by design and op code.....	25
<b>Figure 12.</b> Virtex-5QV DSP RegisterSusceptibility to Protons. Units are cm <sup>2</sup> per register; each DSP has one M register and four others: A, B, C, and P although register use is optional and, therefore, varies by design and op code. ....	26
<b>Fig. 13a.</b> Rotation sweeps for dual-node configuration cells storing ones [Ref. 12, Fig. 23 pt 1] .....	29
<b>Fig. 13b.</b> Rotation sweeps for dual-node configuration cells storing ones [Ref. 12, Fig. 23 pt 2] .....	30
<b>Fig. 14a.</b> Rotation sweeps for dual-node configuration cells storing zeros [Ref.12, Fig.28 pt.1] .....	31
<b>Fig. 14b.</b> Rotation sweeps for dual-node configuration cells storing zeros [Ref.12, Fig.28 pt.2] .....	32
<b>Figure 15.</b> Virtex-5QV Configuration Bit Upset Susceptibility due to Protons .....	35

## TABLE OF TABLES

<b>Table 1.</b> Comparison of Architecture Resources.....	2
<b>Table 2.</b> Latchup Test DUT Conditions .....	3
<b>Table 3.</b> Latchup Test Data .....	5
<b>Table 4.</b> Weibull Fit Parameters for Virtex-5QV SEFIs .....	15
<b>Table 5.</b> Virtex-5QV Heavy-Ion Weibull Fit Parameters for Static Tests .....	17
<b>Table 6.</b> Virtex-5QV Proton Weibull Fit Parameters for Static Tests.....	17
<b>Table 7.</b> Virtex-5QV Heavy-Ion BRAM HICUP Model Fit Parameters .....	20
<b>Table 8.</b> CREME96 HUP and PUP Parameters for Static SEE Rate Calculation.....	27
<b>Table 9.</b> Virtex-5QV Heavy-Ion Weibull Fit Parameters for an Equivalently Hard, Single-Node Configuration Cell.....	36

# 1            **VIRTEX-5 OVERVIEW**

The Xilinx Virtex-5 device is a static random access memory (SRAM)-based, in-system, reconfigurable field programmable gate array (FPGA). The Virtex-5 architecture includes seven major, programmable block types optimized for specific functions:

- The Configurable Logic Blocks (CLB) provide functional elements for combinatorial and synchronous logic, including configurable storage elements, cascadable arithmetic functions and, newly introduced in Virtex-5, 6-input lookup tables (LUTs).
- The Digital Signal Processing (DSP) Slices provide advanced high-speed arithmetic and comparison functions, including multiply and accumulate.
- The Block Memory modules provide large 36-Kbit storage elements of true dual port RAM.
- The Digital Clock Manager (DCM) blocks provide clock frequency synthesis and de-skew.
- The Phase Lock Loop (PLL) blocks provide clock distribution delay compensation, clock multiplication/division, coarse/fine-grained clock phase shifting, and input clock jitter filtering.
- The bidirectional Input/Output Blocks (IOB) have optional Single Data Rate (SDR) or Double Data Rate (DDR) registers and serializers and deserializers (SERDES) enabling support for many industry input/output (I/O) standards, plus selectable drive strengths and digitally controlled output impedance.
- High speed serial GTX transceivers are capable of running up to 4.25 Gb/s. Each GTX transceiver supports full-duplex, clock-and-data recovery.

The high-reliability, radiation-hard Virtex-5QV product is a family of one part type, the XQR5VFX130. Its counterpart in the commercial Virtex-5 line is the XC5VFX130T and they share exactly the same type and numbers of basic blocks and capabilities except for the removal of the System Monitor and PowerPC blocks and a few I/O's in the corners (including two GTX transceivers) in order to improve package reliability. The rad-hard version incorporates several upset-hard-by-design elements, including:

- Dual-node configuration cells that require charge collection in at least two active nodes before an upset can occur.
- Dual-node master-slave user flip-flops (user registers).
- Single-event transient filters on all CLB flip-flop inputs: data, clock, and control.
- Triple-modular redundancy (TMR) in the control circuitry and registers.

Table 1 lists the main architectural resources of the Virtex-5QV space-grade FPGA in comparison with its commercial-grade counterpart. These two parts share a compatible superset footprint to allow prototyping with the commercial device [1]. See the overview datasheet [2] and the electrical specification datasheet [3] for more details on the Virtex-5QV and pointers to more detailed documentation.

**Table 1.** Comparison of Architecture Resources

	Description	Commercial-Grade XC5VFX130T	Space-Grade XQR5VFX130
CFG*	Configuration Bits* (millions)	34.1	34.1
BRAM	Block Memory Bits (millions)	10.9	10.9
LOGIC	Slices (4 Lookup Tables/slice)	20,480	20,480
DSP	18x25 Multiply, 48 bit Accumulate	320	320
PPC	PowerPC405 Processors	2	none
CMT**	Clock Manager Tiles	6	6
MGT	High-speed Transceivers	20	18
IOBs	Input/Output Blocks	840	836

\* Only real memory cells in the Configuration Bit Stream are counted here (not counting BRAM)

\*\* Each CMT includes two Digital Clock Managers (DCM) and one Phase-Locked Loop (PLL)

Like its commercial counterpart, the Virtex-5QV device is fabricated in a 65-nm process geometry. That makes it currently the most highly scaled complementary metal oxide semiconductor (CMOS) technology offered to the aerospace industry. For Xilinx, it is the first product with extensive radiation-hard by design (RHBD) features; Virtex-4QV and earlier space-grade FPGAs use exactly the same mask and circuitry as a particular revision of their commercial counterpart.

This report is the result of the combined efforts of members within the Xilinx Radiation Test Consortium (XRTC), occasionally known as the Xilinx SEE Test Consortium. The XRTC is a voluntary association of aerospace entities, including leading aerospace companies, universities and national laboratories, combining resources to characterize reconfigurable FPGAs for aerospace applications. Previous presentations and publications of Virtex-5QV radiation results have been made by Consortium members, notably at the NSREC and MAPLD conferences and the SEE Symposium.

This report focuses on the measured upset characteristics of the main static (or unlocked) memory elements. A companion report is planned that will cover individual functional blocks including frequency effects. Reports similar to this one documenting the upset susceptibility to heavy ions and protons of the static memory elements in the older space-grade Virtex-4QV [4] and Virtex-2 [5] families are available; Virtex-4QV is still recommended as a viable choice for new designs. A comparison with those results shows that the upset-hard-by-design features of the Virtex-5QV are remarkably successful, making even more of a giant leap forward in improved upset characteristics than the enviable improvements in speed, density and architectural features derived from and accompanying the process scaling to 65 nm.



## 2 LATCHUP TESTING

Fully tested production Virtex-5QV devices were tested for Single-Event Effect (SEE)-induced latchup events at the Texas A&M Cyclotron Institute in July of 2011. Two prior tests had been run on pre-production samples. There were no beam-induced latchup events recorded during any of those tests.

Table 2 provides a summary of the device under test (DUT) test parameters for the latchup irradiations. Each DUT was heated to a nominal pretest temperature of 120 -125°C and biased with specification-maximum voltages.

**Table 2.** Latchup Test DUT Conditions

Parameter	Value	Unit
DUT junction temperature (target)	+120	°C
Internal voltage	1.05	V
I/O voltage	3.45	V
Auxiliary voltage	2.66	V

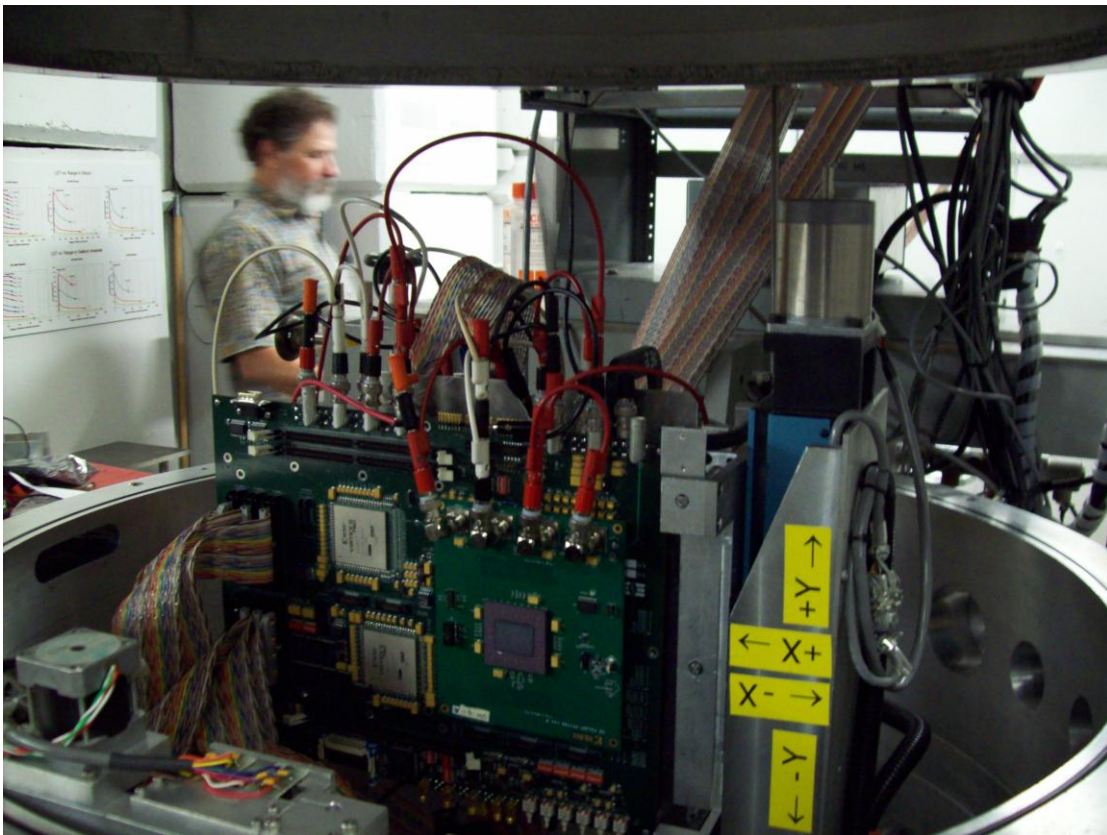
Latchup is the high current state of a bistable parasitic four-layer (PNPN) structure inherent in CMOS where a short circuit sustains itself through positive feedback. It may be triggered by single-event charge deposition or electrical noise, but the only way to remove latchup is to cycle power or , more formally, the structure's voltage must be dropped to a level below the "holding voltage" (preferably before physical damage occurs). For the purpose of these experiments, a practical definition of this classic latchup phenomenon was adopted: any observation of a (1) sudden high current mode that (2) occurred during irradiation that (3) stopped normal operation and (4) required a power cycle of the DUT in order to recover functionality and nominal current would be counted as a latchup event.

During the test runs, the DUT core and I/O voltages and their dynamic currents were captured and recorded in a running log (strip chart). Current trigger levels were set above the maximum nominal draw, so that, in the event of a latchup condition, the monitoring software would detect that the set thresholds were exceeded, declare a latchup condition and quickly cycle power to prevent device damage. In order to reach the desired high fluence in a reasonable period of time, very high fluxes were used for the latchup testing. While satisfying to use obtain a fluence of irradiation equivalent to millions or even billions of years in space, the practical reality is that, even with RHBD cells, the fastest scrubbing is not fast enough to maintain functionality of the implemented design at these highly accelerated irradiation rates. Further, due to the high target fluence for these experiments, it was expected that the DUT would be subject to multiple Single Event Functional Interrupt (SEFI) conditions during the run but at a rate where they could be counted accurately (this turned out to be the case and the SEFI results are in line with those coming in Section 3.2). Demonstrating full recovery of post-beam design operation without the need for a device power cycle proves that no latchup occurred in the irradiation. Therefore, the test procedure adopted was as follows:

1. Program and readback to verify DUT configuration memory; verify design functionality.

2. Heat DUT to +120°C.
3. Record initial temperature, voltage, and current conditions.
4. While irradiating the DUT to a fluence of at least  $10^7$  particles/cm<sup>2</sup>, record DUT voltages and currents via power supplies and the internal diode temperature via dedicated readout chip attached directly to the DUT's pins.
5. Scrub and readback DUT configuration memory after the end of irradiation and verify design functionality.

The tests were performed in vacuum with backside-thinned samples. Care was taken to ensure that ion range was enough so that the ions would actually exit the silicon before they reached their Bragg peak energy deposition; thus end-of-range effects and their accompanying uncertainties and complications were avoided. Photo 1 shows the test motherboard and a DUT card mounted in the vacuum chamber.



**Photo 1.** Test Apparatus Setup in Vacuum Chamber at Texas A&M

Table 3 shows the run parameters and results for each DUT. A 15 MeV/amu Au ion beam with an effective Linear Energy Transfer (LET)  $> 104$  MeV/mg/cm<sup>2</sup> was used; a fluence greater than  $2 \cdot 10^7$  particles/cm<sup>2</sup> was deposited in each run. In some cases, multiple test runs were conducted and added in order to obtain the total fluences shown in Table 3. Tilting the samples with respect to the beam was employed to increase the effective LET to the values shown. The effective residual range listed is the thickness of another (hypothetical) layer of silicon that the ions could penetrate after exiting the real physical silicon; in fact, they penetrate into the dead layers beyond (metal and insulator).

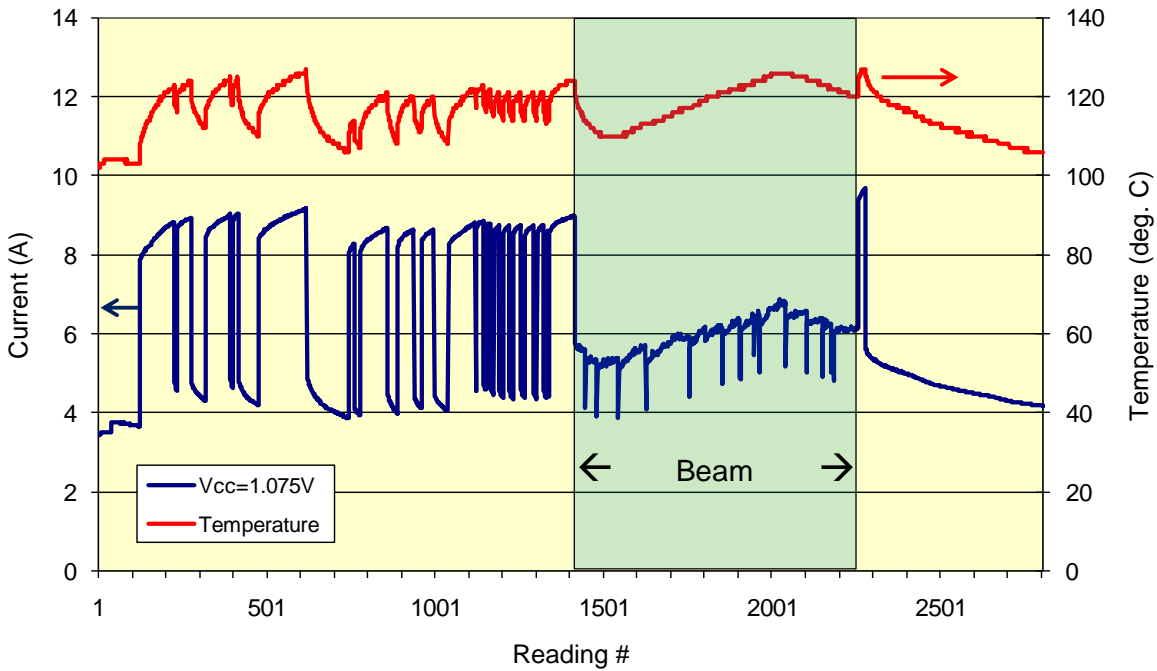
**Table 3. Latchup Test Data**

Using the Texas A&amp;M Cyclotron's 15 MeV/amu Gold (Au) Beam

<u>XQR5VFX130</u> DUT serial #	Effective LET, MeV-cm <sup>2</sup> /mg	Effective Range, μm	Average Flux, #/cm <sup>2</sup> -s	Total Fluence, #/cm <sup>2</sup>	Max Temp °C	Min Temp °C	Single -Event Latchups
s/n:515	104.8	65.2	1.8x10 <sup>5</sup>	8.0x10 <sup>7</sup>	125	117	none
s/n:515	132.2	40.6	2.8x10 <sup>5</sup>	2.0x10 <sup>7</sup>	126	115	none
s/n:515	145.2	30.6	3.1x10 <sup>5</sup>	2.0x10 <sup>7</sup>	127	120	none
s/n:514	145.5	35.6	1.9x10 <sup>5</sup>	2.0x10 <sup>7</sup>	126	114	none
s/n:592	135.7	67.5	0.90x10 <sup>5</sup>	2.0x10 <sup>7</sup>	126	110	none

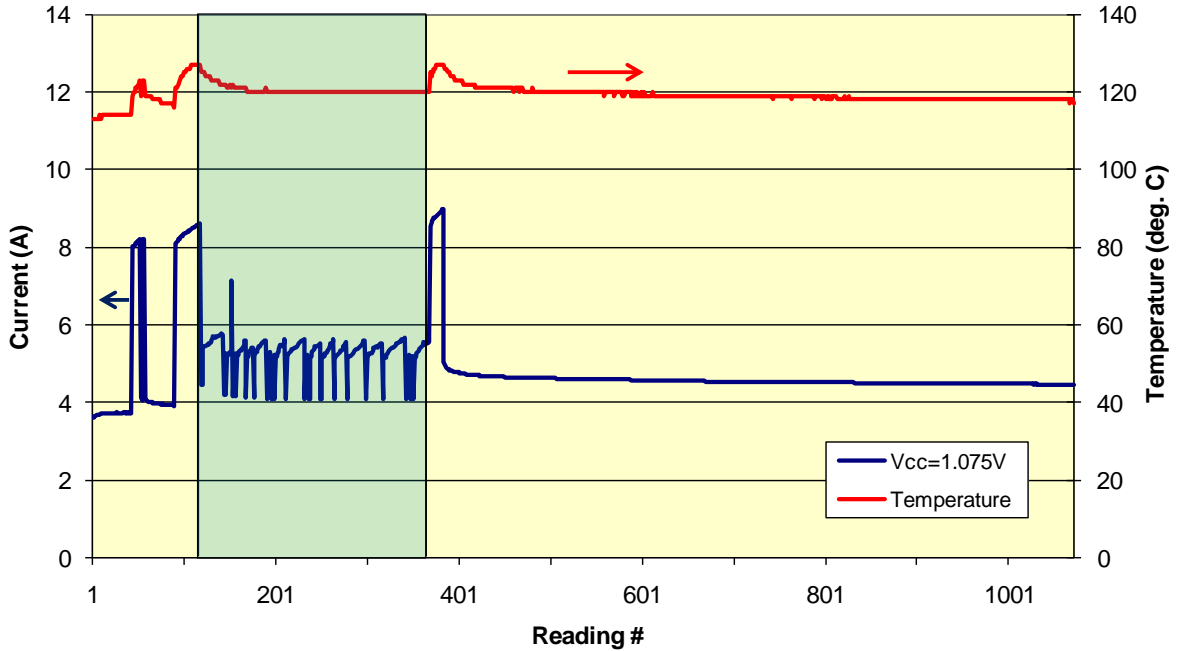
Because Virtex-5 devices are only offered in flip-chip packaging, irradiation is done through the backside of the silicon substrate. In order to reach the active layer at the bottom with sufficiently high-LET, the backside of the silicon must be thinned to less than 100 μm. Figures 1a-1c show the recorded current and temperature before, during, and after the irradiations for the last three rows of Table 3.

Because the bottom of the silicon is solder “bumped” to a fully populated ball-grid package, it is difficult to heat the device with an external heating element to the desired temperature for latchup testing. In order to obtain junction temperatures near 125°C, devices were configured with a “heater” design meant mainly to consume dynamic current and cause self-heating. The core temperature of the device is monitored by measuring the resistance of an internal diode specifically provided for this purpose and controlled by gating the design’s clock.



**Figure 1a.** Current and temperature for s/n: 592 during latchup testing: angle=50 degrees, fluence =  $2.0 \times 10^7$  gold ions per sq.cm, effective LET=135.7 MeV per mg/cm<sup>2</sup>.

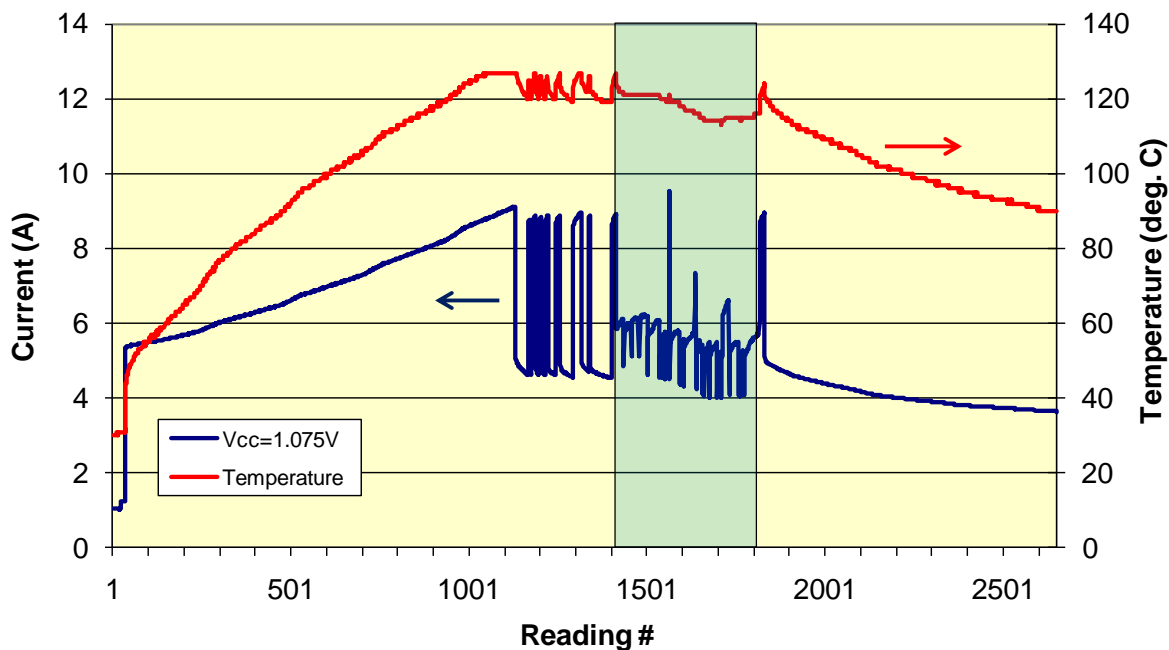
The heater design is a long shift-register chain of CLB flip-flops, long enough to consume more than 90% of the available device resources. The start of the register chain is fed by a one-bit counter so that alternating ones and zeros advance through the chain with each clock pulse. A digital clock management block or DCM multiplies the input clock frequency and, thus, increases the self-heating effect. The clock on/off is manually controlled by the experimenter to increase or maintain temperature pre-beam. The roughly 2A vertical current changes in the strip charts of Figure 1 outside of the irradiation periods correspond to the experimenter exercising that manual control. When the clock is disabled the static only current is drawn and the diode temperature immediately drops as radiative and conduction cooling paths begin to dominate. When the clock is enabled, approximately two amps of dynamic current is added to the current draw and the diode temperature rises as a result.



**Figure 1b.** Current and temperature for s/n:515 during latchup testing: angle=50 degrees, fluence =  $2.0 \times 10^7$  gold ions per sq.cm, effective LET=145.2 MeV per mg/cm<sup>2</sup>

Early in each irradiation, the configuration of the device was upset and the dynamic contribution to self heating would stop. In spite of the high fluxes used which, in turn, led to a high rate of upset, occasionally scrubbing was able to restore functionality and sometimes those were seen by the power supply and appear as positive current spikes during the irradiation in Figure 1b and 1c. Note that negative spikes correspond to SEFI events (detailed in Section 3.2); these were automatically detected and cleared via subsequent reconfiguration from pulsing the PROG\_B input. Supplemental heating was provided under the DUT on the daughterboard using ohmic flat strip heaters driven by a manually controlled power supply. In Figure 1b, two adjustments of the strip heating are shown clearly: first, an increase early in the run when the temperature had dropped to 110°C near the beginning of irradiation and, the second, a decrease near the end when the

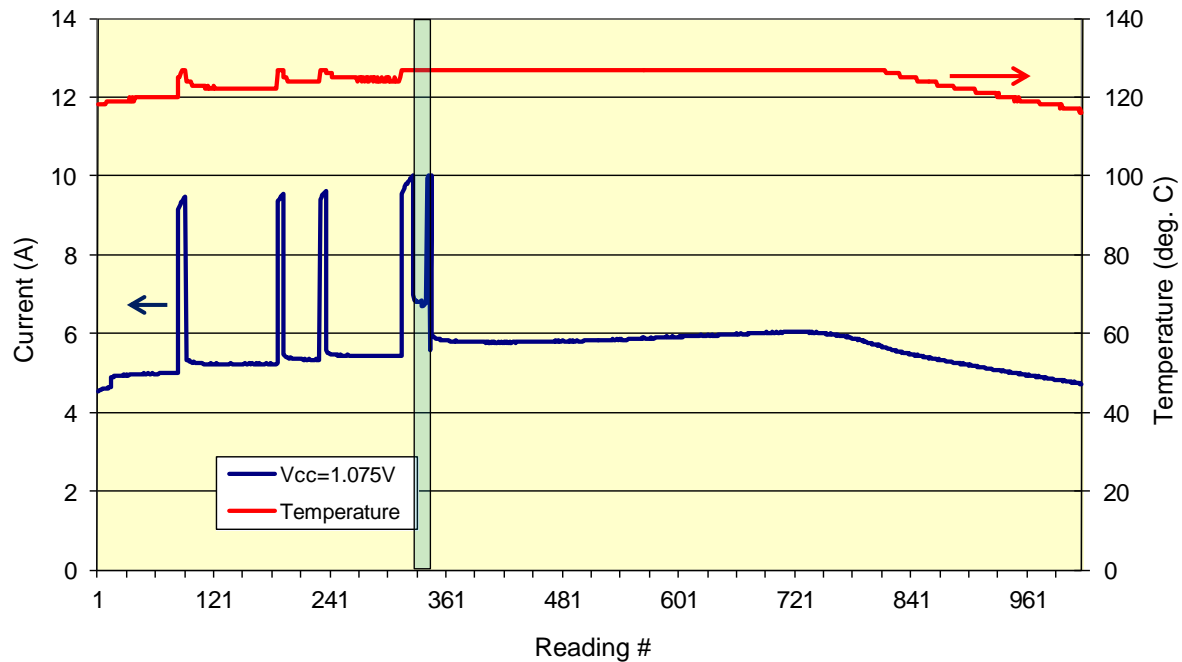
temperature reached 126°C. The instrumentation used had one degree resolution and maxed out at 127°C, that is, an indicated 127°C could have been higher. Maximum and minimum temperatures over the course the irradiations are noted in Table 3.



**Figure 1c.** Current and temperature for s/n:514 during latchup testing: angle=50 degrees, fluence =  $2.0 \times 10^7$  gold ions per sq.cm, effective LET=145.5 MeV per mg/cm<sup>2</sup>

It should be noted that the power supply software reported one high current reading just before the facility beam shutter opened and one suspected latchup event just after the shutter was closed for the first irradiation of s/n:592 (before that shown in Figure 1a). Initially, the experimenters and facility personnel suspected that this was the result of a latchup induced by shutter electromagnetic noise pulse (EMP) rather than the beam itself. Recent beam shutter replacements had left out some of the usual noise-suppression capacitor/resistor circuitry and, indeed, some shutter-coincident anomalous upsets were seen by some XRTC experiments- and others- in June 2011. However, a detailed examination of the logs revealed that the DUT, in fact, did not experience latchup. The test software current threshold setting was exceeded due to an overheating mistake. As can be seen in Figure 2, the temperature was pegged at 127°C; the measured current never exceeded 10.005 A and the voltage never drooped due to current limiting above that as would have been the case for a latchup event. With no clock for 170 seconds after the  $\approx 6.7$  second irradiation (the experimenter's reaction time plus shutter lag), the temperature reading finally came back into range indicating that heating lag had significantly overheated the DUT and that the high current reading seen by the software indicated correct operation for the out-of-spec high temperature, and was not latchup. To put it another way, the current draw of nominal operation at the high temperature some distance above the 127°C instrumentation limit was higher than the software latchup threshold setting. Buttressing this conclusion, the Functional Monitor and Configuration Monitor logs (shown at the top of Figure 3) recorded that the DUT was operating normally and

correctly after the beam shutter opened for a few seconds. At that point, the software cycled the power supply after two consecutive readings of 10.005A (actually correct operation), slightly in excess of the selected threshold of 10.000A. Note that the following irradiation of the same DUT (s/n:592) using identical conditions, except for more careful temperature control, is the one shown back in Figure 1a where the software detects no latchup through a fluence of  $2.0 \times 10^7$  ions/cm<sup>2</sup>.



**Figure 2.** Current and temperature for s/n: 592 during latchup testing: angle=50 degrees, fluence  $< 2.8 \times 10^5$  gold ions per sq.cm, effective LET=135.7 MeV per mg/cm<sup>2</sup>

In summary, we conclude that no single-event latchup was observed during this quite extreme testing at high voltage, high temperature, high LET and with high fluence.



### 3 STATIC TESTING

The static Single-Event-Upset (SEU) experiments' goal is to measure the upset susceptibility of memory elements incorporated in XQR5VFX130 FPGA for both heavy ions and protons in order to enable calculation of expected upset rates in given space environments. For the dual-node based cells (configuration memory and user flip-flops), a special test chip (dubbed TC-65nm) was required because grazing angle ion hits have significant impact on the space upset rate calculations. Other static characterizations were carried out on the pre-production FPGA samples (dubbed FX-1) and, subsequently, on production FPGA samples. In addition, as an adjunct to all FPGA static, dynamic, and mitigation testing, each DUT was carefully monitored for SEFIs during all experiments. For all these tests, the same basic experimental setup was used: the XRTC motherboard with an appropriate daughter DUT board. Note that the Configuration Monitor capability was not used when testing the TC-65 DUTs as they are more like mini-ASICs and do not have a programmable configuration.

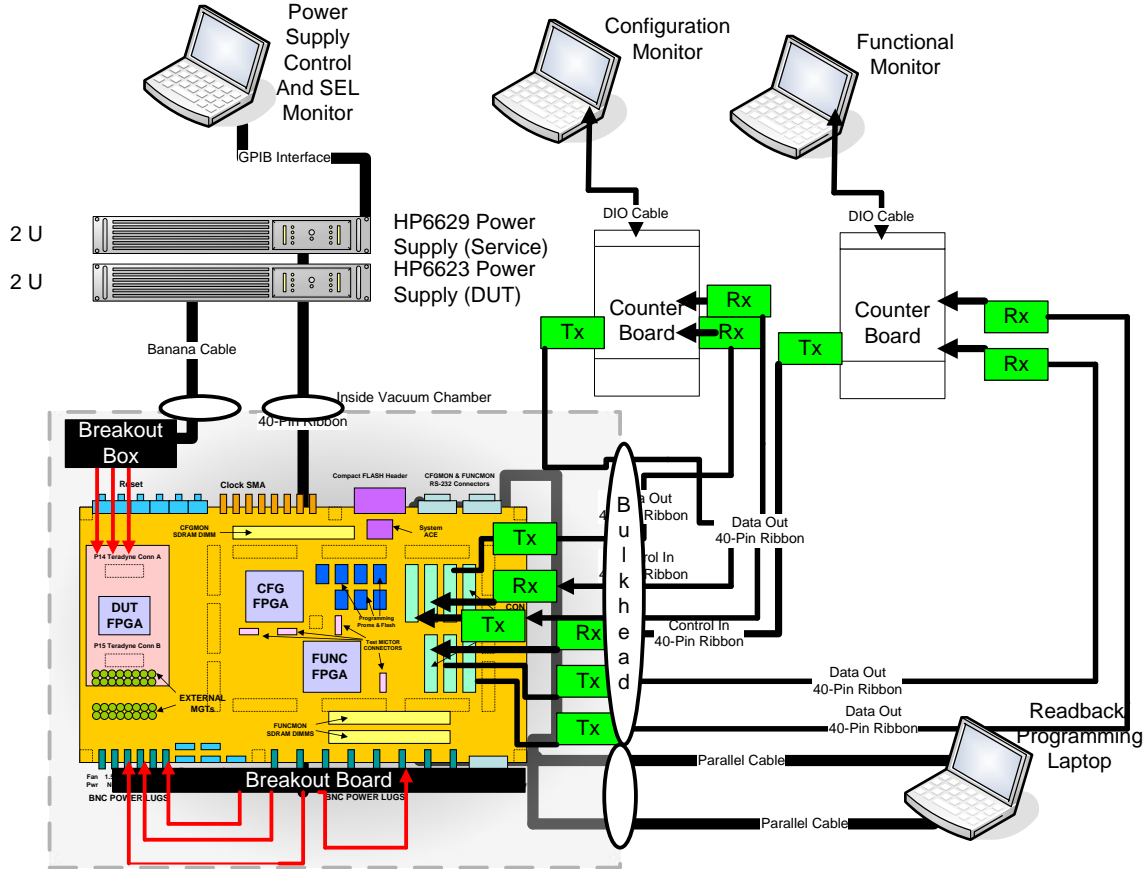
It is important to understand that, as used here, the term "static testing" is mainly meant phenomenologically. This static testing was done without clocking the configured design in the DUT during an irradiation or clocking it at frequencies that are proven to have little effect on the test results, that is, at slow speed. Physically, this does not mean that all upsets are direct upsets of the memory cell under test. Some (or even most) of the upsets may be due to Single-Event Transients (SETs) on clock lines or asynchronous control signals, especially on the upset-hardened-by-design elements where direct upsets of the memory cells were intentionally suppressed.

#### 3.1 Experiment Setup

Figure 3 shows the test setup in vacuum as needed for the latchup testing. This is essentially the same apparatus that supported the earlier Virtex-4QV testing [4]. JPL's 5-40 pin bulkhead was used to run five of the six communication cables through the vacuum chamber. The sixth was run through the 50 pin D-Sub connector provided by the irradiation facility at the Texas A&M University (TAM) cyclotron. Three parallel cables were also sent through the 50 pin connectors (one for a DUT readback Parallel-IV cable, one for a motherboard/DUT design programming Parallel-IV cable, and one for the temperature monitoring circuit).

A mounting platform with integrated power breakout cables was used for mounting the motherboard to the rotating chassis in the vacuum chamber and for extracting the four power supplies from the 40-pin cable. The four supplies were sent through the vacuum chamber bulkhead over BNC connectors then re-integrated to the 40-pin cable. 'Force' and 'Sense' wires were tied together at the power supply (HP6629) for all four supplies and provided the necessary 2.5V, 3.3V, and 3.3V I/O voltages for the motherboard; the last supply was used to control heater strips attached to the back of the daughter card. When used, the receiver/driver cards were powered by the 3.3V of the motherboard I/O although they are often omitted. The 5V supply for the Parallel-IV JTAG cables (replaced by USB programming pods when not running in the vacuum chamber) and temperature sensor circuit were powered by an external Agilent E3610A through the BNC bulkhead. The DUT power supply was an HP6623, which provided three supplies with different hardware current limits of 5A, 10A, and 2A; supply one provided 2.5V for DUT  $V_{AUX}$ , supply two

provided 1.0V for DUT  $V_{INT}$  and supply three provided 3.3V (or, optionally, 2.5V) to the  $V_{CCO}$  DUT I/O banks that talk to the motherboard. ‘Force’ and ‘Sense’ wires were tied together at the bulkheads on supplies one and three which were run through BNC bulkhead feedthroughs. High current cables were used to run ‘Force’ for supply two in through a 40-pin cable bulkhead connector and were adapted back into banana cables with a second custom 40-pin cable inside the chamber; 20 pins were used for power and 20 for ground in the feedthrough. ‘Sense’ for supply two was sent through a BNC bulkhead adapter to banana cables and connected to force at the daughter card (as near to the DUT as possible).



**Figure 3.** SEU Test Hardware Setup

More conveniently, most of the SEU testing was conducted in air with long range ions and extra thin DUTs to ensure penetration. Ion LETs were adjusted for the intervening materials including the air gap and exit window. The setup for in-air testing was essentially the same as in-vacuum testing except, of course, the adapted connections for getting through the bulkheads were eliminated. Also for the in-air irradiations, USB2-based programming pods with USB cable extenders or hubs were used instead of the parallel port programming pods indicated in Figure 3. Because it was shown that reliable communication does not require them, the receiver/driver pairs were not used for most of the testing. Finally, for proton testing, it proved necessary and expedient to avoid interfering secondary neutron upsets to the two monitoring FPGAs so the motherboard was upgraded, chiefly by replacing the old workhorse Virtex-II Pro devices with new production RHBD Virtex-5QV FPGAs.



## 3.2 SEFI Results

Single-event functionality interrupts are an important class of upset phenomena because their effects are big, usually resulting in the need for a complete reset or re-boot of the affected system with an accompanying period of outage of system function. In this section, *device* SEFIs, or those put in by the manufacturer, are the focus rather than *design* SEFIs, or those put in by the system designer. The line can be a bit blurry because how a *design* uses a *device* can expand or contract the universe of possible SEFIs. Ignoring known best practices will increase SEFIs in both type and number. These results do not try to attempt a sensitivity study of those possibilities, but rather to characterize the unavoidable minimum set associated with known best practice.

Traditionally, to recover an FPGA-based system from a device SEFI, the FPGA must be re-configured via pulsing the PROG pin or cycling power and involves a minimum outage of some tens or hundreds of milliseconds. However, results are included here for a new category of potentially intrusive SEFIs that can be reset and functionality recovered without reconfiguration, potentially significantly reducing the malfunction timescale.

### 3.2.1 Experiment Considerations

The heavy ion irradiation test campaign on the Virtex-5QV FPGA was performed at two facilities- the BASE Facility at Lawrence Berkeley National Laboratory (LBL) and The Cyclotron Institute, Texas A&M University (TAM). The first test trip was in November 2008 and the most recent was in September 2012. The static, dynamic, and mitigation tests performed represent the equivalent of many millions of years in the space radiation environment. In these testing, the DUT was always monitored for device SEFIs. Thus, every irradiation was two simultaneous experiments- the specific function under study as well as a SEFI test.

The Virtex-5QV has been tested with various ions at different incident angles covering a LET range of 0.11–145.5 MeV-cm<sup>2</sup>/mg. A combination of beam energy degraders and DUT angles were used to achieve the higher LETs. Note that, in some cases, the effective LET from an angled DUT gave a different cross section than the same LET from a heavier ion at normal incidence; this disagreement seems clear at the 65-nm node and will likely grow as scaling continues. This can be, at least partially, attributed to bigger multiple bit upsets (MBUs) from steeper angles. Because of the aim of using this data is to predict space rates, normal incidence data was strongly preferred.

The proton SEFI test campaign was undertaken starting in November 2011 through December 2012 at the University of California at Davis (UCD) cyclotron. A preliminary test in June 2010 showed that, because the Virtex-5QV shows extremely low proton sensitivity for SEFIs (and also for configuration upsets, see Section 4), the upset rate from secondary neutrons of the supporting Configuration Monitor and Functional Monitor FPGAs was too large for reliable measurements. The proton test campaign was conducted only after a major rework of the XRTC motherboard to replace these FPGAs with Virtex-5QV devices. It was then possible to get the needed measurements but still required a significant cost in terms of beam hours used and parts dosed.

The data plots in this report have two sigma (or 95% confidence limit) statistical error bars on the cross section measurements based on the assumption that the underlying population distribution is Poisson; in some cases, the error bars are smaller than the plotting symbol and, thus, don't show in the plots. When no events were seen, only the top of the error bar appears. Event counts and beam fluence were summed across individual irradiation runs to obtain the cross section reported for a given set of experimental conditions.

Typically, device SEFIs are low probability and are rarely seen on orbit; there are no known cases where a Xilinx FPGA experienced a device SEFI in space. Nevertheless, Xilinx designers have significantly reduced the Virtex-5QV's susceptibility to device SEFIs and hence it is not expected to ever be seen on an actual space mission. However, some SEFIs were observed in the beam testing where event rates are hugely accelerated in order to obtain statistically significant and accurate measurements of low probability events. The original criterion for a SEFI is that it requires either a complete reconfiguration or a power cycle before the device recovers normal operation. In early tests, a very rare SEFI mode was observed that seemed to require power-cycling, but it is now understood that it can be eliminated by following the recommendations in XMP120[6] as the reference design in XAPP588[7] does. More recently, in the proton tests, three of these power-cycle-required events were encountered but subsequent investigation revealed that some of the recommended steps were not in place and, while it is impossible to prove for certain, it is likely that these SEFIs would not have occurred (or would have manifested as ordinary POR SEFIs that respond to the PROG pin) if the recommendations had been followed. This experience should serve as a reminder of the importance of following the recommendations documented in XMP120. Avoiding unnecessary complexity is always desirable and a simpler alternative is to *always* power cycle on SEFI.

### 3.2.2 SEFI Nomenclature and Taxonomy

A great deal of effort has been undertaken to categorize SEFIs into understandable phenomenological categories and to recognize their signatures in key status registers. The knowledge gained is incorporated into these results as well as recommendations for on-orbit configuration management and SEFI detection, if needed [7]. The results below demonstrate that, the probability of an on-orbit SEFI is very low; thus, only missions with the most stringent reliability requirements will need to be concerned with them

The observed SEFIs, for the Virtex-5QV, are placed into two main categories:

1. Design-intrusive
2. Visibility-intrusive

In the design-intrusive category are the Power-On-Reset-like (POR) and the Global Signal (GSIG) SEFIs. The visibility-intrusive category includes the malfunctions of the SelectMap Port (SMAP) or the Frame Address Register (FAR) where upsets may cripple its ability to auto-increment correctly or disable writing it.

The Start-Up (SU) SEFI has proven to be a uniquely less intrusive SEFI. Issuing a startup command restores the DONE pin and the design operation continues undisturbed. However, communication via IOs is interrupted between the onset of the SU SEFI and the receipt of the start-up command; this is, at least, potentially a system problem, perhaps

even requiring re-boot if the system cannot tolerate a brief communication outage. As a result, the SU represents another category of device SEFI:

### 3. Communication-intrusive.

Because it is potentially intrusive, a conservative choice is to treat every DONE drop as an intrusive SEFI and reconfigure; this was the choice made initially and indeed for most of the test campaign. However, reconfiguration is not necessary and recovery via the startup command may be relatively fast. This allows the possibility, particularly attractive for proton-rich environments, that designers may choose to build a system that operates through an SU SEFI. Configuration monitoring for the proton testing was changed to implement this approach, and the tests were naturally robust to the SU SEFI communication outages, even difficult to notice in most cases. Note that the SU SEFI is the only identified SEFI that does not require reconfiguration for recovery.

The Readback (or RB) SEFI experienced by the Virtex-4 [4] where some bits cannot be scrubbed without unmasking them falls into the not-really-a-SEFI category as it is more a scrubbing annoyance than a design or visibility intrusion. The SCRUB SEFI seen in the Virtex-4 [4] was not observed in Virtex-5QV. Most of the testing was done with the Xilinx recommended frame-based scrubbing where checking for an SMAP SEFI between each frame prevents accidentally putting in more than a frame's worth of bad or misaligned data. In theory, it is still possible for this SEFI to occur so following the frame-based scrubbing recommendation is prudent.

The POR SEFI results in a global reset of all internal storage cells and the loss of all program and state data. Observation of this mode is that when it occurs it is almost always accompanied by the DONE pin dropping low, a sudden change of the DUT current to its starting value and the loss of all configured functions. If a configuration readback is attempted, then an unusually large readback error count will be seen (millions of bits in error).

The SMAP SEFI is the loss of either correct read or write capabilities through the SelectMAP port. This SEFI is indicated either by the retrieval of only meaningless data or inability to refresh data. It is possible that in some cases, the port could be re-activated by commanding the JTAG port to find and correct errors in the control registers. However, some SMAP SEFIs also affect the JTAG port as well, making this an incomplete solution. Therefore, a complete reconfiguration was used in these tests to regain full port access and function.

The FAR SEFI exhibits a loss of control of the frame address register, such as it continuously increments uncontrollably. It is detected by an inability to write and read control values to the FAR while all other aspects of the SelectMAP port are still fully functional. For the purpose of orbital error rate calculations, the FAR SEFI is considered a sub-set of SMAP SEFI modes. However, for characterization purposes, it is individually scrutinized, but occurred extremely rarely during heavy ion testing and did not occur in the proton tests.

The Global Signal SEFI was separated from other design-disrupting SEFIs similar to what was done in the earlier Virtex-4 testing [4]. These signals include GSR (Global Set/Reset), GWE\_B (Global Write Enable), GHIGH\_B (Global Drive High), and others. They can all be observed through the status (STAT) register or one of the control (CTLx) registers. The

ones that can be scrubbed fall into the potentially intrusive category along with the SU SEFI, but the others that require a reconfiguration are clearly design intrusive. Note that no Global Signal SEFIs were seen in any of the proton tests.

Readback upsets that can't be scrubbed (no longer categorized as a SEFI) occur when a masked portion of the readback data has been upset and, thus, cannot be corrected. This stuck-in-upset condition is caused by the use of the GLUTMASK, which is needed to enable the use of SRL16s in conjunction with partial reconfiguration. If GLUTMASK is not invoked then un-scrub-able bits are almost completely eliminated. These bits, which cannot be corrected through partial reconfiguration, do not affect the operability of the configured design. In Virtex-5QV, configuration bits that can be accessed through DRP bits fall into this category even when the resource they are associated with is not implemented. In that case, the upset of an un-scrub-able bit does not affect design operation. The exceptional case is when the DRP bit is associated with a resource used by the design. Note that only clocking resources (DCMs and PLLs) and the MGTs have DRP bits; they are there in order to make it possible for a design to change key parameters, like frequency for example, on the fly; thus, they can only be scrubbed directly from the design side, not through a configuration port.

Scrub SEFIs were observed for the first time in the Virtex-4 testing [4], but they were not observed in the extensive Virtex-5QV test campaign. This extremely rare SEFI mode is the result of an upset causing corruption of the data stream being scrubbed into the DUT. This obviously can disrupt the design operation and may be accompanied by some large internal contention currents. On the Virtex-4 a key characteristic of the Scrub SEFI is that it is the only SEFI that has ever been observed to have a design dependence, that is, it exhibit a different cross sections for different designs. It may be prudent to follow the Virtex-4 recommendation of checking for SEFIs frequently to limit the extent of any Scrub SEFI that may occur. Frame-based scrubbing where checks are done between each frame was used in most of the test campaign and that fact partially explains why no Scrub SEFIs were encountered.

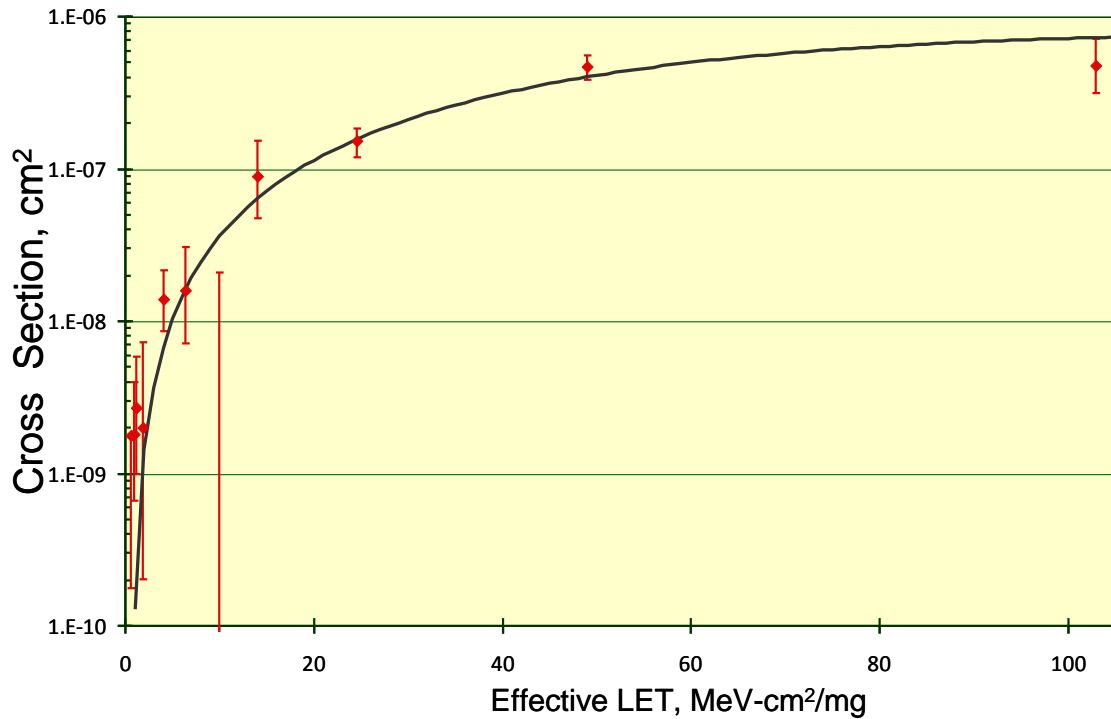
### *3.2.3 SEFI Measurements and Data Fits*

SEFI experimental results are given below. The fit parameters of the Weibull curves are given in Table 4 and the curves and data shown in Figures 4a and 4b. In the heavy ion results of Figure 4a, POR and SMAP SEFIs are, by far, the most prominent components. While the exact proportions vary with LET, overall the design-intrusive SEFIs are about 70% while visibility SEFIs are about 20%. The remaining 10% are communication type events known as SU SEFIs that can be recovered via re-issuing a startup command rather than having to reconfigure.

**Table 4.** Weibull Fit Parameters for Virtex-5QV SEFIs

SEFI	Weibull Parameters			
	Limit (cm <sup>2</sup> /devi ce)	Onset -	Width -	Power -
Heavy-Ion Combined	$8.0 \times 10^{-7}$	0.7 MeV-cm <sup>2</sup> /mg	60	1.65
Proton Combined	$1.1 \times 10^{-13}$	5 MeV	10	0.8
Proton. SUs	$8.0 \times 10^{-14}$	5 MeV	5	2

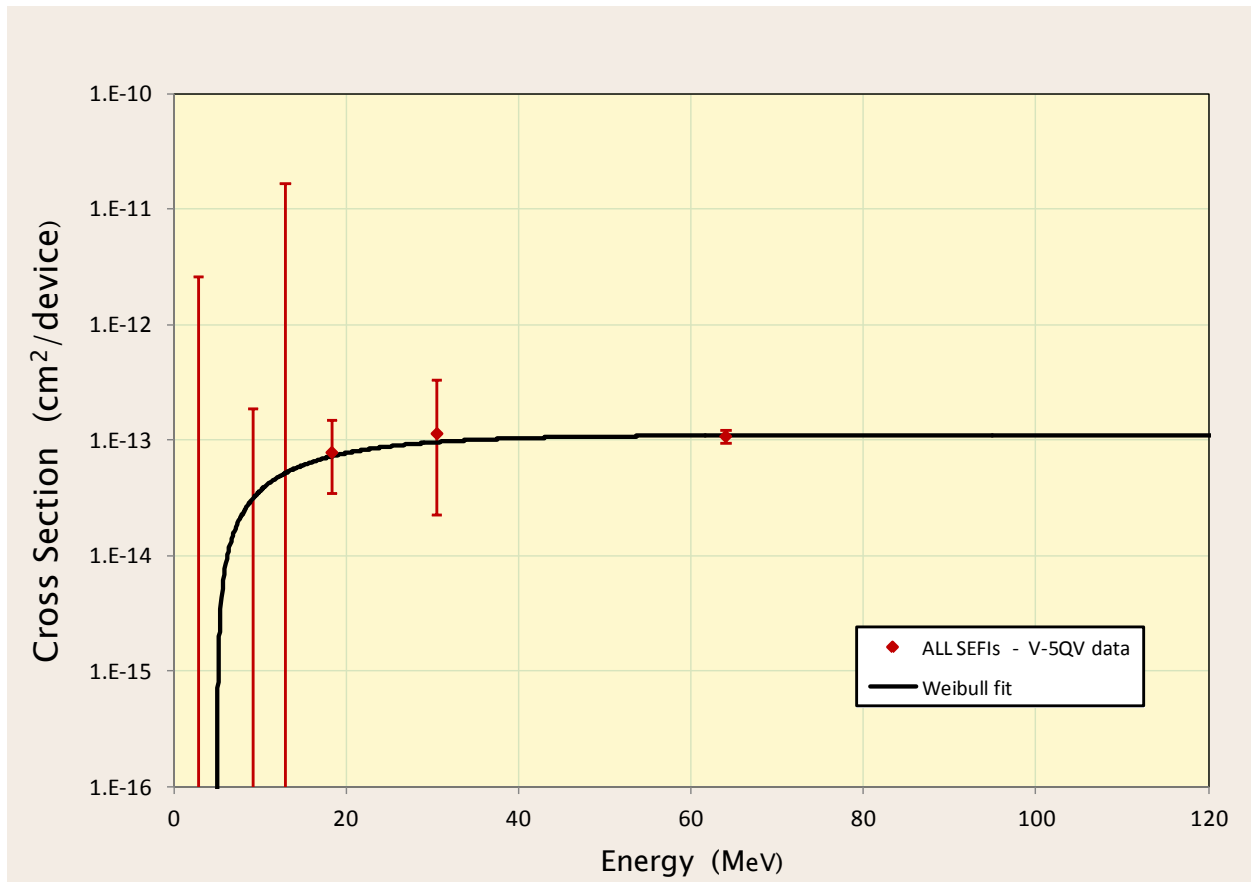
During the proton tests, only POR, SMAP, and SU SEFIs were encountered with SU SEFIs clearly the most likely. A portion of each of these is mitigated via built-in TMR in the configuration state machine and, therefore, the measured cross sections for that portion will be strongly dependent on the beam flux used (essentially quadratic). To extrapolate TMR-ed circuitry from beam rates to space rates requires good data at two or three widely spaced fluxes. Acquiring such data is not really a tractable task. So we are left with conservative, probably quite conservative, measurements to use as the basis of space rate predictions. The higher fluxes used by necessity in the both the proton and heavy ion testing are will cause orders of magnitude more TMR-mitigated SEFIs for a given fluence than the same fluence at the much lower irradiation rate of the real space environment.

**Figure 4a.** Virtex-5QV Susceptibility to ALL SEFIs from Heavy Ion Strikes

With regards to the clearly design intrusive categories, only the POR SEFI was encountered in proton testing. Because the on-orbit proton rates are not affected very much by the shape of the “knee” of the curve at low energy, resources were concentrated

on getting reasonable statistical significance for the POR SEFI at 64 MeV only. Using a step-function like fit to this high energy data point does add a bit of conservatism into the rate predictions. That was deemed a reasonable tradeoff instead of expending many tens of beam hours and dozens of parts needed for statistically significant cross section points at lower energy. The upset-hard cross sections are just too low, especially at low proton energy.

Along with the change to Virtex-5QV service FPGAs, two other procedural changes accompanied the 2011-2012 proton tests: (1) strict monitoring of the Configuration Monitor FPGA itself for upsets and (2) no PROG pin assertion on SU SEFI detection. The first change means that runs were terminated immediately if SEFI monitoring became even potentially untrustworthy. The latter change allowed recovery from an SU SEFI via issuing the startup command to be observed and the effect, if any, recorded in the Functional Monitor log. Close inspection of the logs was required for most test types in order to detect the brief communication SEFI's existence; only the IO-specific tests had non-subtle signatures. Because SU SEFI type was dominant in the proton testing, a separate fit for the SU SEFIs is included in Table 4. Note that the SU SEFIs are included in the "ALL" data and fits.



**Figure 4b.** Virtex-5QV SEFI Susceptibility due to Protons

### 3.3 Other Static Results

This section presents results for three distinct memory element types included in the Virtex-5QV- Block RAM, User Registers, and DSP Registers. Summarizing the results, Weibull parameters for heavy ion fits to the data are given in Table 5. Table 6 is a placeholder for the proton data fits. Data plots for each memory element are given in the following sub-sections while additional details are given in the referenced papers and reports.

**Table 5.** Virtex-5QV Heavy-Ion Weibull Fit Parameters for Static Tests

Memory Type	Weibull Parameters			
	Limit, cm <sup>2</sup> /bit	Onset, MeV-cm <sup>2</sup> /mg	Width -	Power -
BRAM	$1.15 \times 10^{-7}$	0.01	200	0.86
User F/F's, Filter=OFF	$2.80 \times 10^{-8}$	0.50	20	2.0
User F/F's, Filter=ON	$2.76 \times 10^{-9}$	0.89	26	2.2
DSP, M register	$5.5 \times 10^{-6}$	0.1	67	1.12
DSP, other registers	$2.0 \times 10^{-6}$	0.1	35	1.25

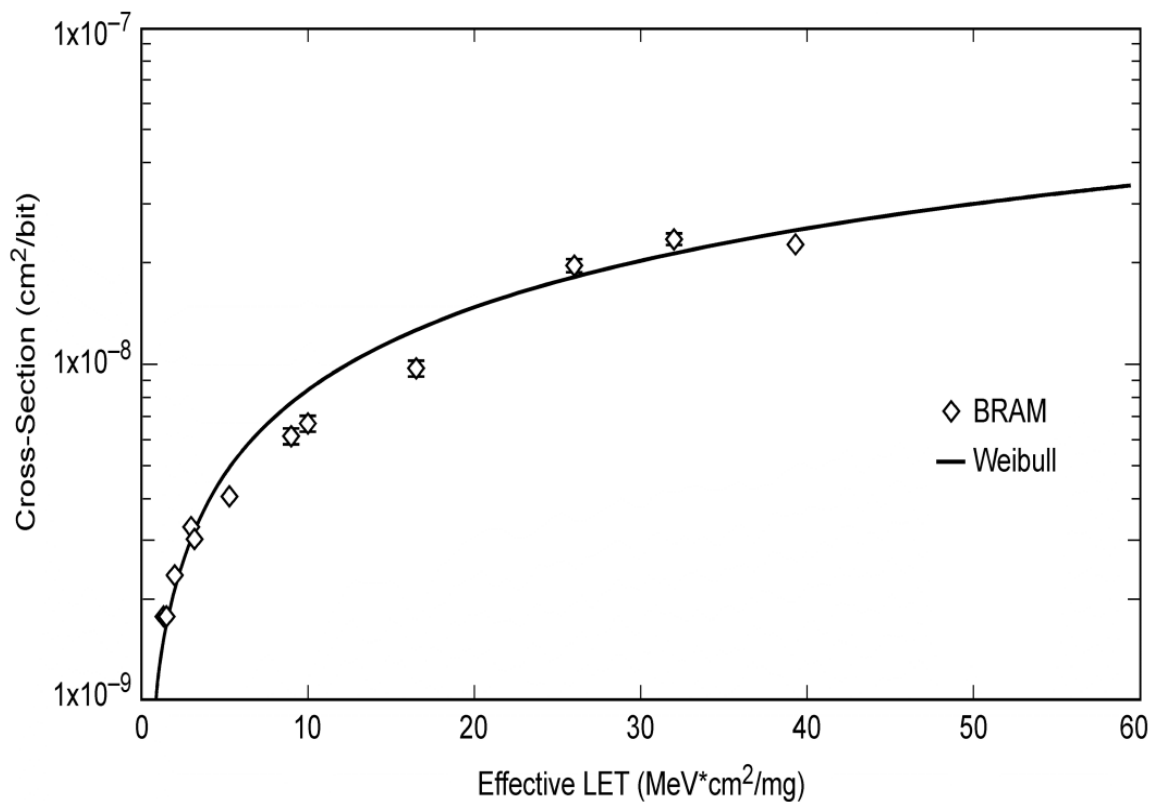
**Table 6.** Virtex-5QV Proton Weibull Fit Parameters for Static Tests

Memory Type	Weibull Parameters			
	Limit, cm <sup>2</sup> /bit	Onset, MeV	Width -	Power -
BRAM	$4.7 \times 10^{-14}$	0.8	12	0.6
User F/F's, Filter=OFF	$3.0 \times 10^{-15}$	5	1	1.0
User F/F's, Filter=ON*	$2.5 \times 10^{-16}$	5	1	1.0
DSP, M register	$9.5 \times 10^{-10}$	5.0	16	2.0
DSP, other registers	$5.3 \times 10^{-10}$	5.7	16	1.4

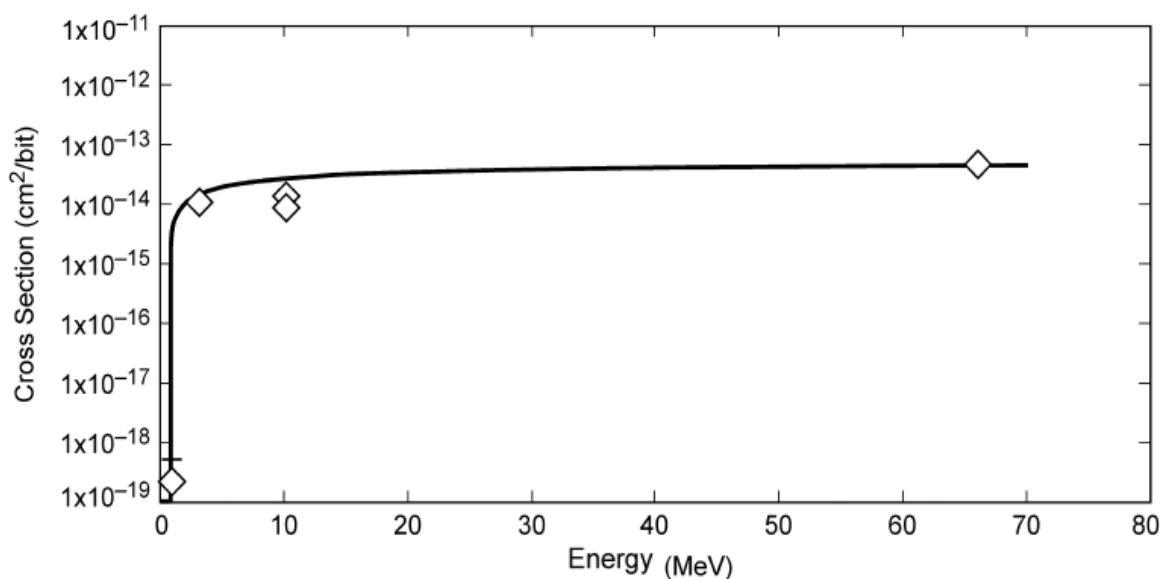
\* - no measured upsets, top of 95% error bar curve, may be overly conservative

### 3.3.1 Block Memory

Normal incidence static heavy ion and proton cross section curves are given in the paper published in 2011 [8] and reproduced in Figures 5 and 6. The fit parameters were given above in Tables 5 and 6.



**Figure 5.** Virtex-5QV BRAM Susceptibility due to Heavy Ion Strikes



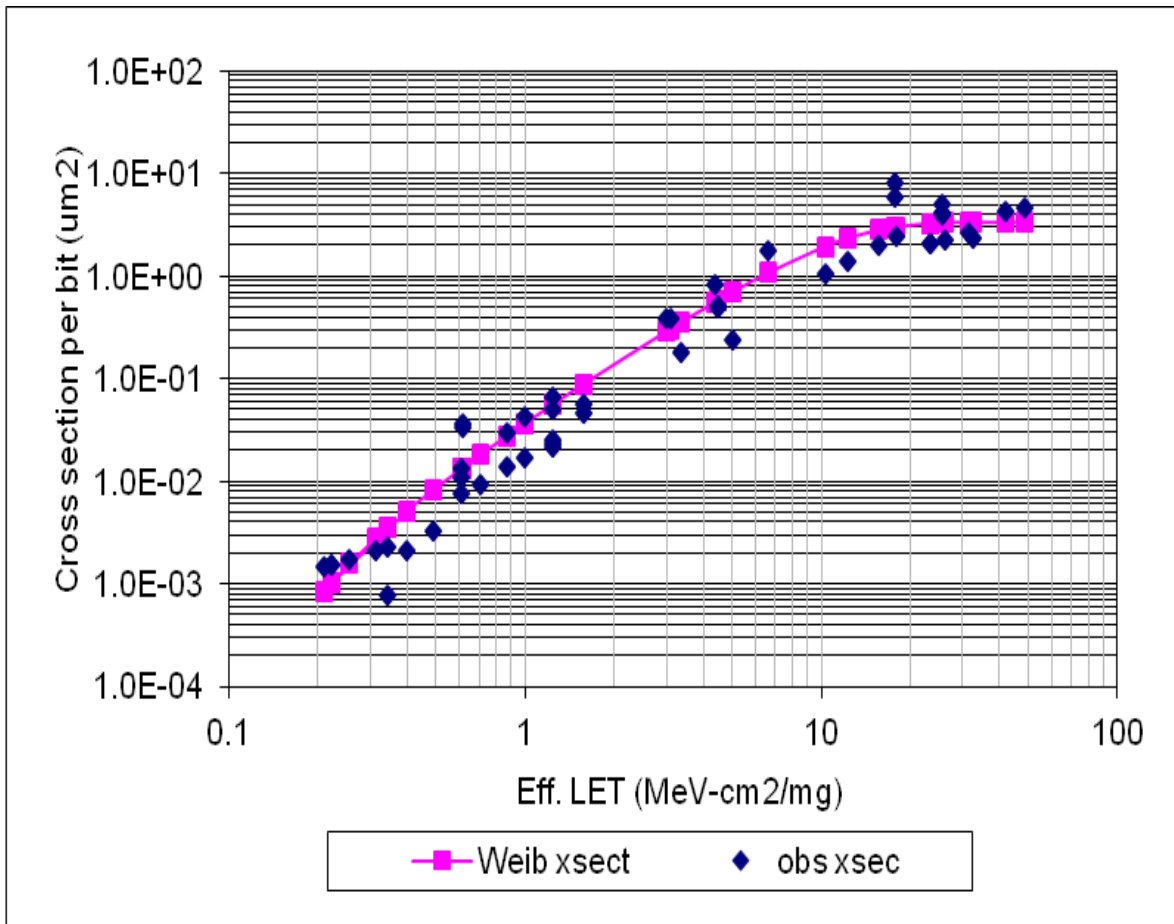
**Figure 6.** Virtex-5QV BRAM Susceptibility due to Proton Strikes



### MBUs –

The above data is taken at or near normal incidence. In July 2011, an XRTC experiment was performed at the Teaxs A&M cyclotron under the direction of Munir Shoga using steeper DUT angles in order to understand the BRAM multiple-bit upset (MBU) susceptibility. It was expected that physical interleaving of the bits would make the error detection and correction circuitry robust to MBUs and that turned out to be the case.

Nevertheless, the use of tilt angles in this experiment yielded a much broader coverage of effective LET than the earlier experiment. There were 50 irradiation runs covering a range of effective LET from 0.2 – 48 MeV-cm<sup>2</sup>/mg at a several angles ranging from 0-75°. The test data and a fit to the data are shown in Figure 7. The data fitting method used here is based on the HICUP model described in Ref. 9 and other publications referenced therein.



**Figure 7.** Virtex-5QV BRAM Susceptibility Angle Experiment v. HICUP Model Fit

It is standard practice to model the sensitive (or charge-collection) volume as a rectangular parallel piped (RPP) where all charge deposited in the volume appears at the circuit node and none deposited outside. The standard model does not really specify how to determine the charge collection depth and sometimes a funnel length is added which enlarges the

sensitive volume. As a result, the RPP dimensions, esp. depth, do not necessarily reflect physical reality, but are fitting parameters. In addition, charge sharing (what if adjacent RPPs overlap?) is not well handled in the “standard” model. The HICUP model attempts to overcome these shortcomings; it is still based on the RPP geometry, but allows a rectangular, that is, non-square, top face, makes some allowance for ion track structure, and uses the angular-dependent heavy ion upset cross section in a more physical way. Seven parameters are adjusted to attain a best fit to the data. These are the dimensions of the RPP (length, width, and height) and the Weibull distribution parameters: the limiting cross section, the onset threshold LET, the width (W), and the power (P). These parameters are shown in Table 7 and can be used directly in RPP-based rate calculations like CREME96.

**Table 7.** Virtex-5QV Heavy-Ion BRAM HICUP Model Fit Parameters

	<b>Weibull Parameters</b>	
Limit	$3.33 \times 10^{-8}$	cm <sup>2</sup> /bit
Onset	0.10	MeV-cm <sup>2</sup> /mg
Width	10.86	MeV-cm <sup>2</sup> /mg
Power	1.75	-
	<b>RPP Parameters</b>	
Length	1.29	μ
Width	2.58	μ
Height	1.70	μ

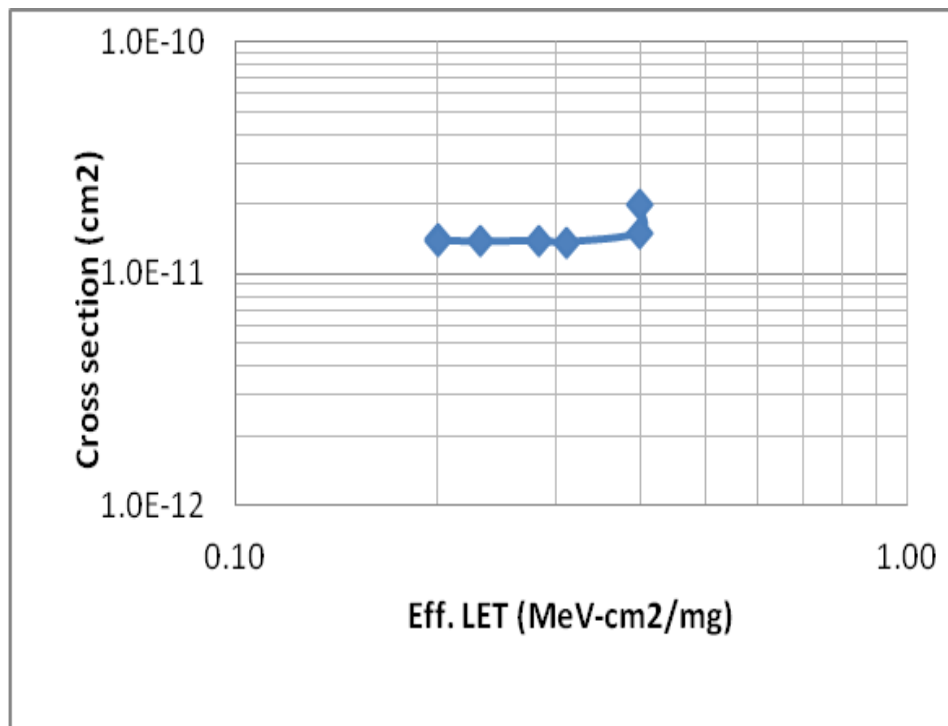
Note that the HICUP model fit parameters of Table 7 give a GEO rate about a factor of five lower than the Weibull parameters of Table 5. This is mainly explained by the addition of the low LET measurements (below 2.0) of the later experiment where most of the rate is coming from; the parameters of Table 5 over-extrapolate the cross section in this critical region. There also should be some smaller effect from underlying differences between the implied model behind standard practice and the HICUP model. The recommended parameters in Table 8 are based on those of the HICUP model; the older parameters given in Table 5 yield overly conservative (that is, too high) on-orbit rates.

#### Direct Proton Ionization Upsets –

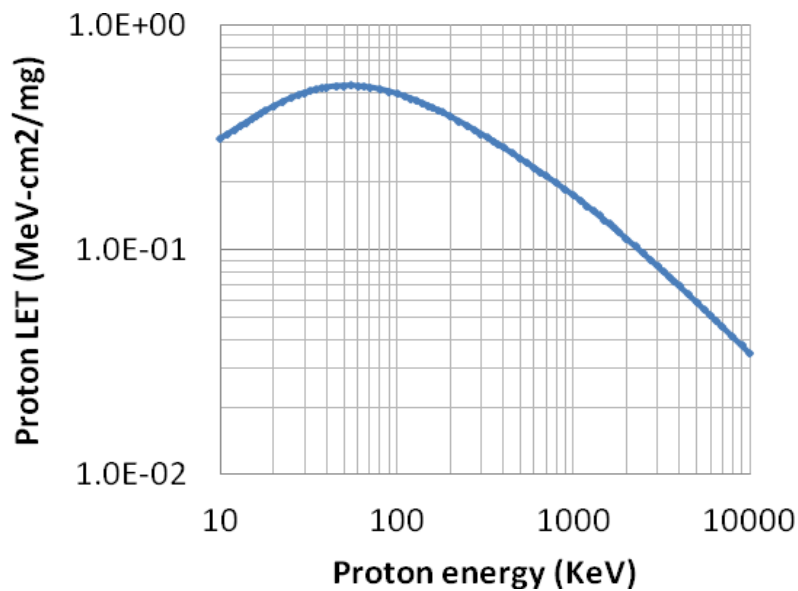
Testing was also done to determine BRAM susceptibility to low energy protons (< 1MeV). Since testing with low energy protons is challenging, testing was performed with He and N ions. Testing with these ions provides more confidence and accuracy in the measured LET for upset. Figure 8 shows a plot of the low energy direct ionization upset cross section vs. LET. The upset cross section is  $\approx 1.5 \times 10^{-11}$  cm<sup>2</sup>/bit. Testing at lower than LET of 0.2 MeV-cm<sup>2</sup>/mg is possible but the upset cross section is expected to be lower because the maximum upset cross section is expected at the highest LET for protons, which occurs at the Bragg peak  $\approx 0.53$  MeV-cm<sup>2</sup>/mg. Data taken at LETs higher than 0.53 MeV-cm<sup>2</sup>/mg are not shown in Figure 8 because those data are irrelevant to this discussion of *direct* ionization. Of course, in the relatively rare event of a proton-induced reaction, protons can deposit more energy *indirectly* through their reaction products.

For reference, the low energy proton LET spectrum is shown in Figure 9. The Bragg peak, where the maximum LET occurs, is at  $\approx 50$  KeV. Notice that the LET drops below 0.2 MeV-cm<sup>2</sup>/mg at an energy of 800 KeV and lower. By approximately this point, a proton

does not have enough LET or residual energy and so cannot deposit enough charge to upset the BRAM cell. For similar reasons, the relatively constant upset cross section in Figure 8 that occurs at LET range of 0.2 – 0.4 MeV-cm<sup>2</sup>/mg is expected.



**Figure 8.** Upset cross section vs. Effective LET for He Irradiation



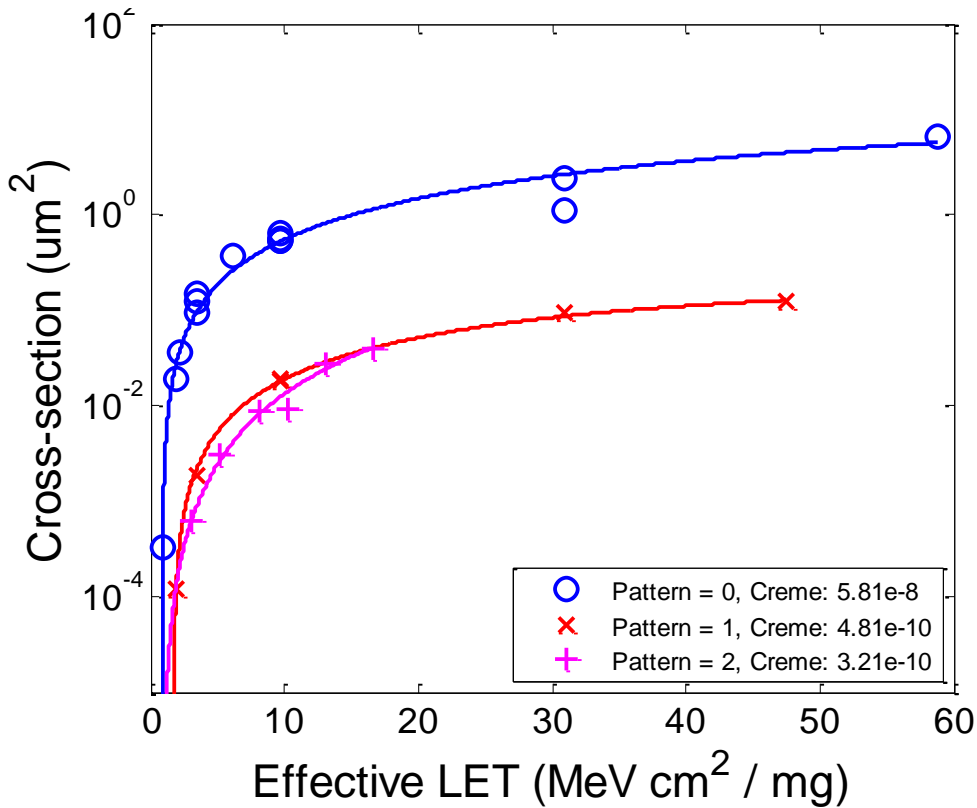
**Figure 9.** Low Energy Proton LET Spectrum

BRAM cells upset at a low Linear LET of  $\approx 0.2 \text{ MeV-cm}^2/\text{mg}$ . This indicates a susceptibility to direct ionization with low energy protons ( $<1\text{MeV}$ ). However, the 4-word interleave separation assures that direct ionization upsets will not cause a DBU within a single ECC word.

From the heavy ion data fit, the saturated per-bit upset cross section is  $\approx 3.3 \text{ } \mu\text{m}^2$ , which is larger than the physical area of a BRAM cell. This is attributed to the effect of track structure and charge transport dynamics, and also, inducing multi bit upset (MBU) with one ion hit.

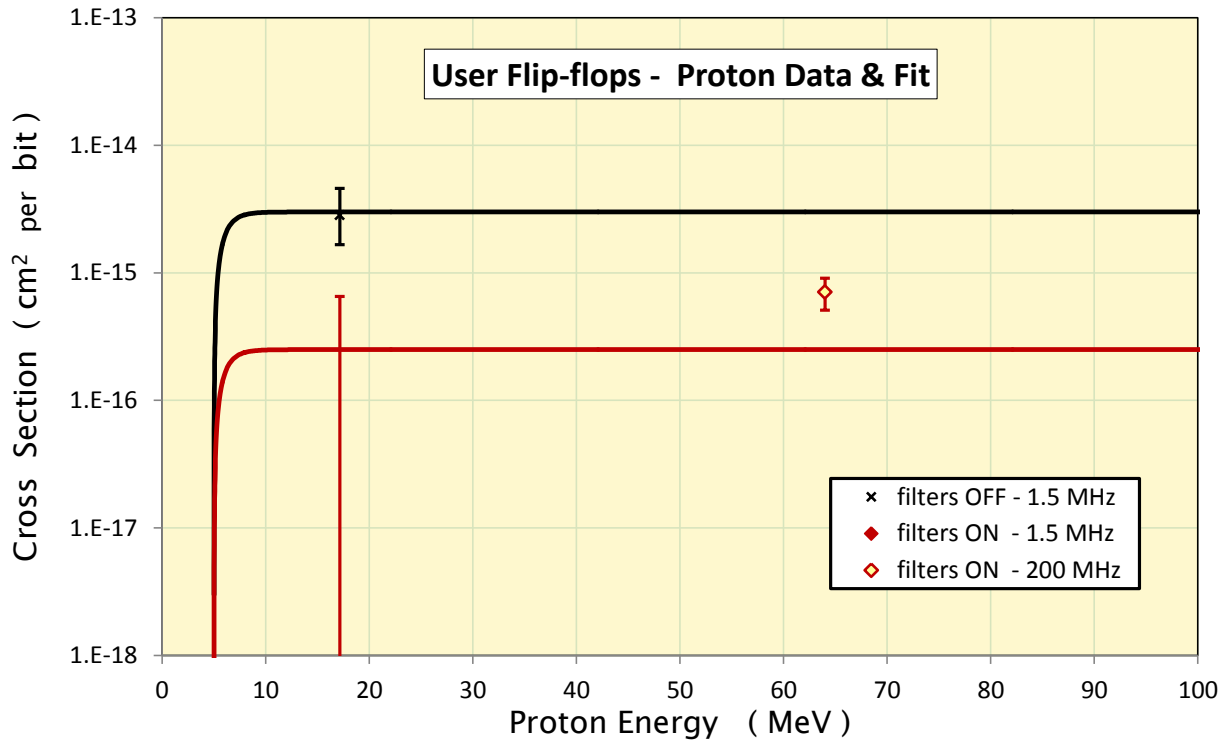
### 3.3.2 User Flip-Flops

Pseudo-static (1.5MHz) data and results are available thanks to George Madias and Eric Miller of Boeing and are summarized in this section. Because the flip-flops are implemented with master-slave dual-node cells, they are very hard to direct hits. Almost all of the upsets observed in the pseudo-static tests are the result of large enough single-event transients on flip-flop inputs, one of: (1) the asynchronous reset, (2) the clock input at a time coincident with inverse data line input, or (3) the data input coincident with the clock edge. As a result, the measured upset susceptibility depends heavily on whether the input SET filters are turned on or not. Also, mainly due to #3, the results are frequency dependent; high speed measurements have been done; those results are the companion dynamic report, *Virtex-5QV Architectural Features SEU Characterization Summary*.



**Figure 10a.** Virtex-5QV User Flip-Flop with filters OFF Pseudo-static (1.5 MHz)  
Susceptibility to Heavy Ions (note: 0=checker, 1=all ones, 2=all zeros)

Proton data collection for the pseudo-static case was purposefully sparse. Almost all proton data were taken operating at 200MHz testing where upsets susceptibilities are higher than the static case. These dynamic results are low enough that the static proton results (especially with the filters ON) will be only of extreme academic interest. For practical purposes, the user flip-flops with filter ON are nearly immune to proton-induced static upsets. The filter-ON fit in Table 6 and Table 8 is based on zero observed upsets and is guided by the top of the 95% (or  $\approx 2$  sigma) error bar for an observed zero at 17.1 MeV. For reference, if one upset had been observed, the measured cross would be  $1.77 \times 10^{-16}$  cm<sup>2</sup> per bit.

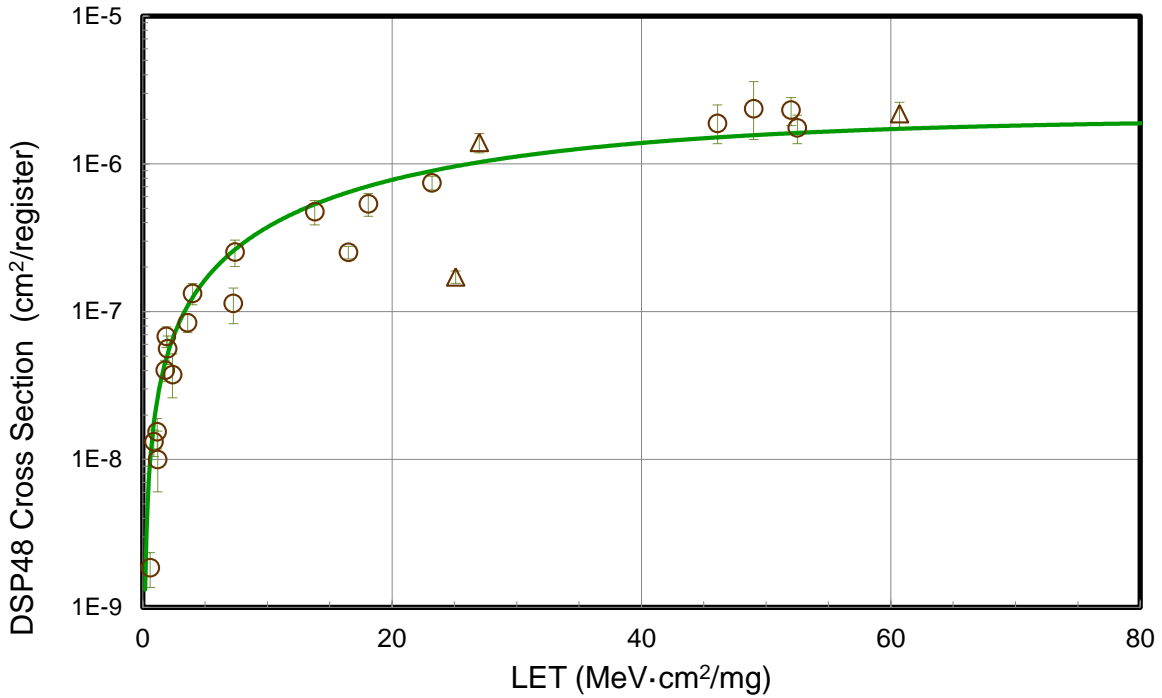


**Figure 10b.** Measured Virtex-5QV User Flip-Flop Pseudo-static (1.5 MHz) Susceptibility to Protons and Weibull Fits. The data is sparse, to say the least, esp. with filters ON so a 64 MeV, 200 MHz “anchor” point is shown.

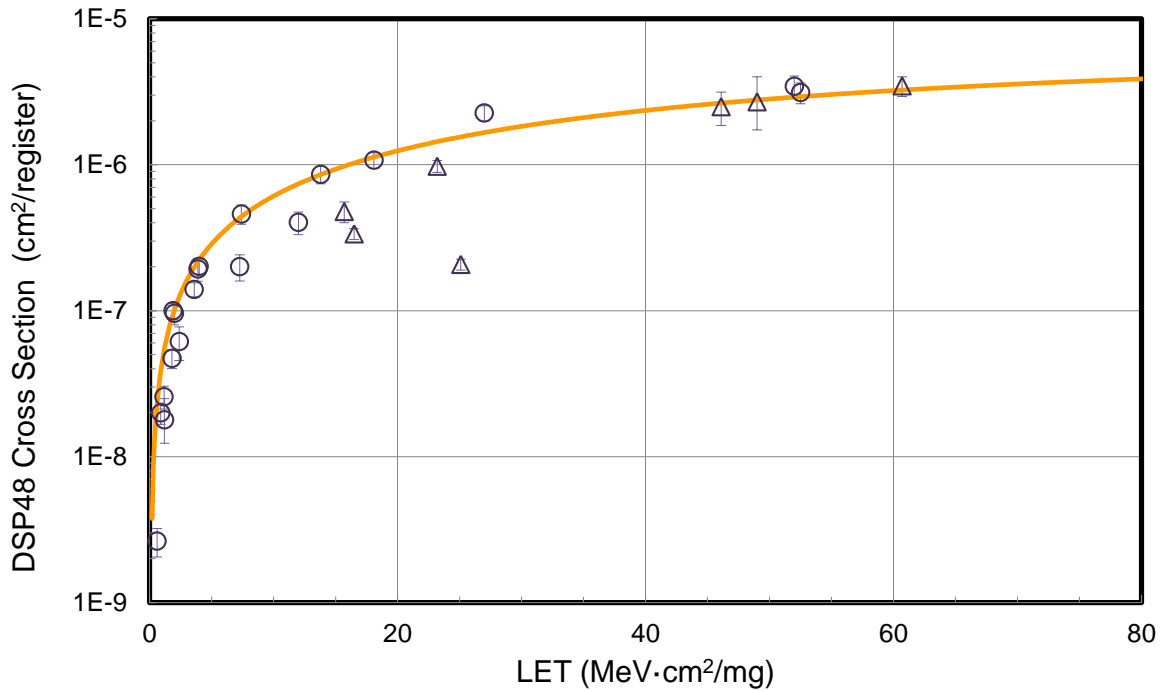
### 3.3.3 DSP Registers

A full report [10] on the XRTC static and dynamic testing of the Digital Signal Processing blocks (DSPs) of the Virtex-5QV by Roberto Monreal of SwRI includes detailed DSP SEU results. This section presents only a brief digest of the static results and includes some newer data.

The DSP registers all appear to have about the same heavy ion static upset susceptibility (see Figure 11a) except for the M register which is somewhat more susceptible so they are shown separately in Figure 11b. Many upset events affect more than one bit in a given register so the cross sections in this digest are all  $\text{cm}^2$  per register. In practice, the bit error rate is of less importance than the register error rate because a calculation uses registers, not bits, as their fundamental granularity; thus, predicting the register upset rate is the more correct way to predict the calculation error rate. Of course, each registers' duty cycle must be folded into the error rate based on the particular op code's usage.

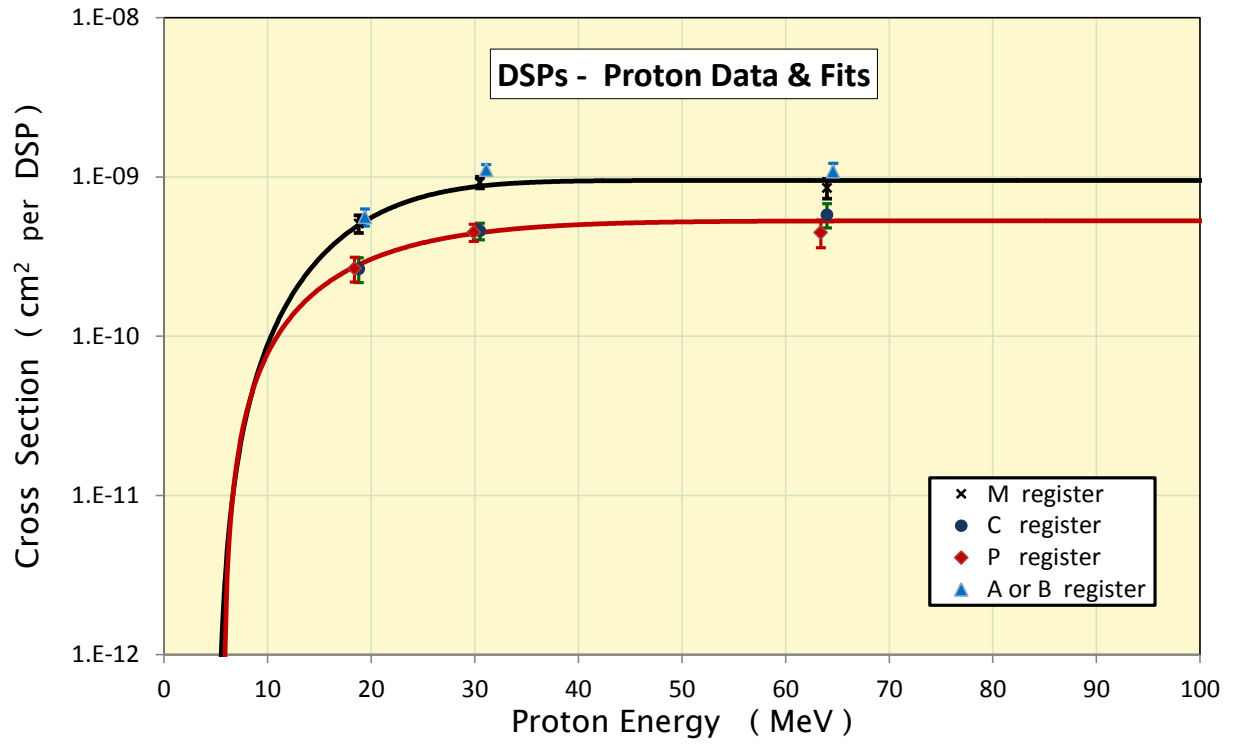


**Figure 11a.** Virtex-5QV DSP A/B-, C-, and P-Register Susceptibility to Heavy Ions; C-Register data shown. Note that each DSP has one M register and three others: A/B, C, and P although register use is optional; therefore, duty cycles vary by design and op code.



**Figure 11b.** Virtex-5QV DSP M-Register Susceptibility to Heavy Ions. Note that each DSP has one M register and three others: A/B, C, and P although register use is optional; therefore, the duty cycle varies by design and op code.

Static proton data on the DSP registers was been taken at the UC-Davis cyclotron in January and June 2012 and the results are shown in Figure 12. Unexpectedly, the A/B data moved up to match the more sensitive M register; while there are some clues that this result is spuriously high and relates to a noise margin problem on the new motherboard, using the M register curve for all DSP registers will not raise the rates very much and, thus, it is a convenient and conservative approach.



**Figure 12.** Virtex-5QV DSP RegisterSusceptibility to Protons. Units are cm<sup>2</sup> per register; each DSP has one M register and four others: A, B, C, and P although register use is optional and, therefore, varies by design and op code.

Heavy ion and proton Weibull fit parameters for the DSP curves are given in Tables 5 and 6, respectively, and also in Table 8.



### 3.4 Orbital Rate Calculations

The CREME96 (Cosmic Ray Effects on Micro-Electronic Circuits) orbital event rate estimation model originally provided by the Naval Research Laboratories and now supported by Vanderbilt University [11] can be used to calculate orbital error rates based on the Weibull fits in the previous sections. Table 8 shows the input parameters for the CREME96 HUP and PUP files for calculating heavy ion and proton induced events, respectively. Although the SEFI data is measured in events per device, for modeling purposes, the bits per device used in CREME96 calculations are adjusted to give a relative per bit cross-section value more typical of a standard register. Even though most SEFI events are caused by logic gate transients, CREME96 models events as static upsets on a storage cell. The same is true for proton upsets of configuration cells (CFGs); what's calculated here are transient-induced as explained in Section 4.3. While determining how many gates can generate those transients is theoretically needed, it is convenient - and doesn't change the results too much anyway - to use the actual CFG bit count.

**Table 8.** CREME96 HUP and PUP Parameters for Static SEE Rate Calculation

<b>CREME96 Input Parameters</b>								
<b>Device</b>	<b>CFG*</b>	<b>BRAM</b>	<b>User F/F</b>		<b>DSP M-Reg</b>	<b>DSP other</b>	<b>Total SEFIs</b>	<b>Units</b>
			<b>f=OFF</b>	<b>f=ON</b>				
<b>FX130</b>	$34.1 \times 10^6$	$10.9 \times 10^6$	81,920		320	1280	4	bits / device**
<b>Sigma(HI)</b>	$1 \times 10^{-8}$ ***	$3.33 \times 10^{-8}$	$2.80 \times 10^{-8}$	$2.76 \times 10^{-7}$	$5.5 \times 10^{-6}$	$2.0 \times 10^{-6}$	$2.0 \times 10^{-7}$	$\text{cm}^2/\text{bit}$ **
<b>Sigma (P)</b>	$2.5 \times 10^{-18}$	$4.7 \times 10^{-14}$	$3.0 \times 10^{-15}$	$2.5 \times 10^{-16}$ †	$9.5 \times 10^{-10}$	$5.3 \times 10^{-10}$	$2.8 \times 10^{-14}$	
<b>Proton (PUP)</b>								
<b>Onset</b>	5	0.8	5	5	5	5.7	5	MeV
<b>Width</b>	50	12	1	1	16	16	20	w
<b>Power</b>	1	0.6	1.0	1.0	2.0	1.4	0.8	s
<b>Limit</b>	0.0000025	0.047	0.003	0.00025	950	530	0.028	$\text{cm}^2/10^{-12}$
<b>Heavy Ion (HUP)</b>								
<b>X &amp; Y</b>	1***	1.82	1.67	0.525	23.5	14.1	4.5	$\mu$
<b>Z + funnel</b>	1***	1	1	1	1	1	1	$\mu$
<b>Onset</b>	0.25***	0.1	0.5	0.89	0.1	0.1	0.7	$\text{MeV}/\text{cm}^2/\text{mg}$
<b>Width</b>	100***	10.86	20	26	67	35	60	w
<b>Power</b>	2.95***	1.75	2.0	2.2	1.12	1.25	1.65	s
<b>Limit</b>	1***	3.33	2.80	0.276	550	200	20	$\mu^2$

\* Not all configuration cells control design elements (so, in those, upsets can't make a design malfunction).

\*\* For DSP registers (not bits): there are 320 of each type in a device and the sigma units are  $\text{cm}^2/\text{register}$

\*\*\* "Equivalent" single-node Weibull for actual dual-node response when storing a one (worst case); see Section 4.4

† No measured upsets, may be overly conservative, construction is similar to configuration cells which are 100x lower

Representative CREME96 orbital error rate estimates for several select orbits will be done for Quiet Solar Minimum conditions and assuming 100 mils aluminum-equivalent spacecraft shielding. All the rates make the conservative assumption that all bits are used; for more accuracy in a given application, scale these results by the fraction of the resources actually used.

## 4 CONFIGURATION (DUAL-NODE) RESULTS

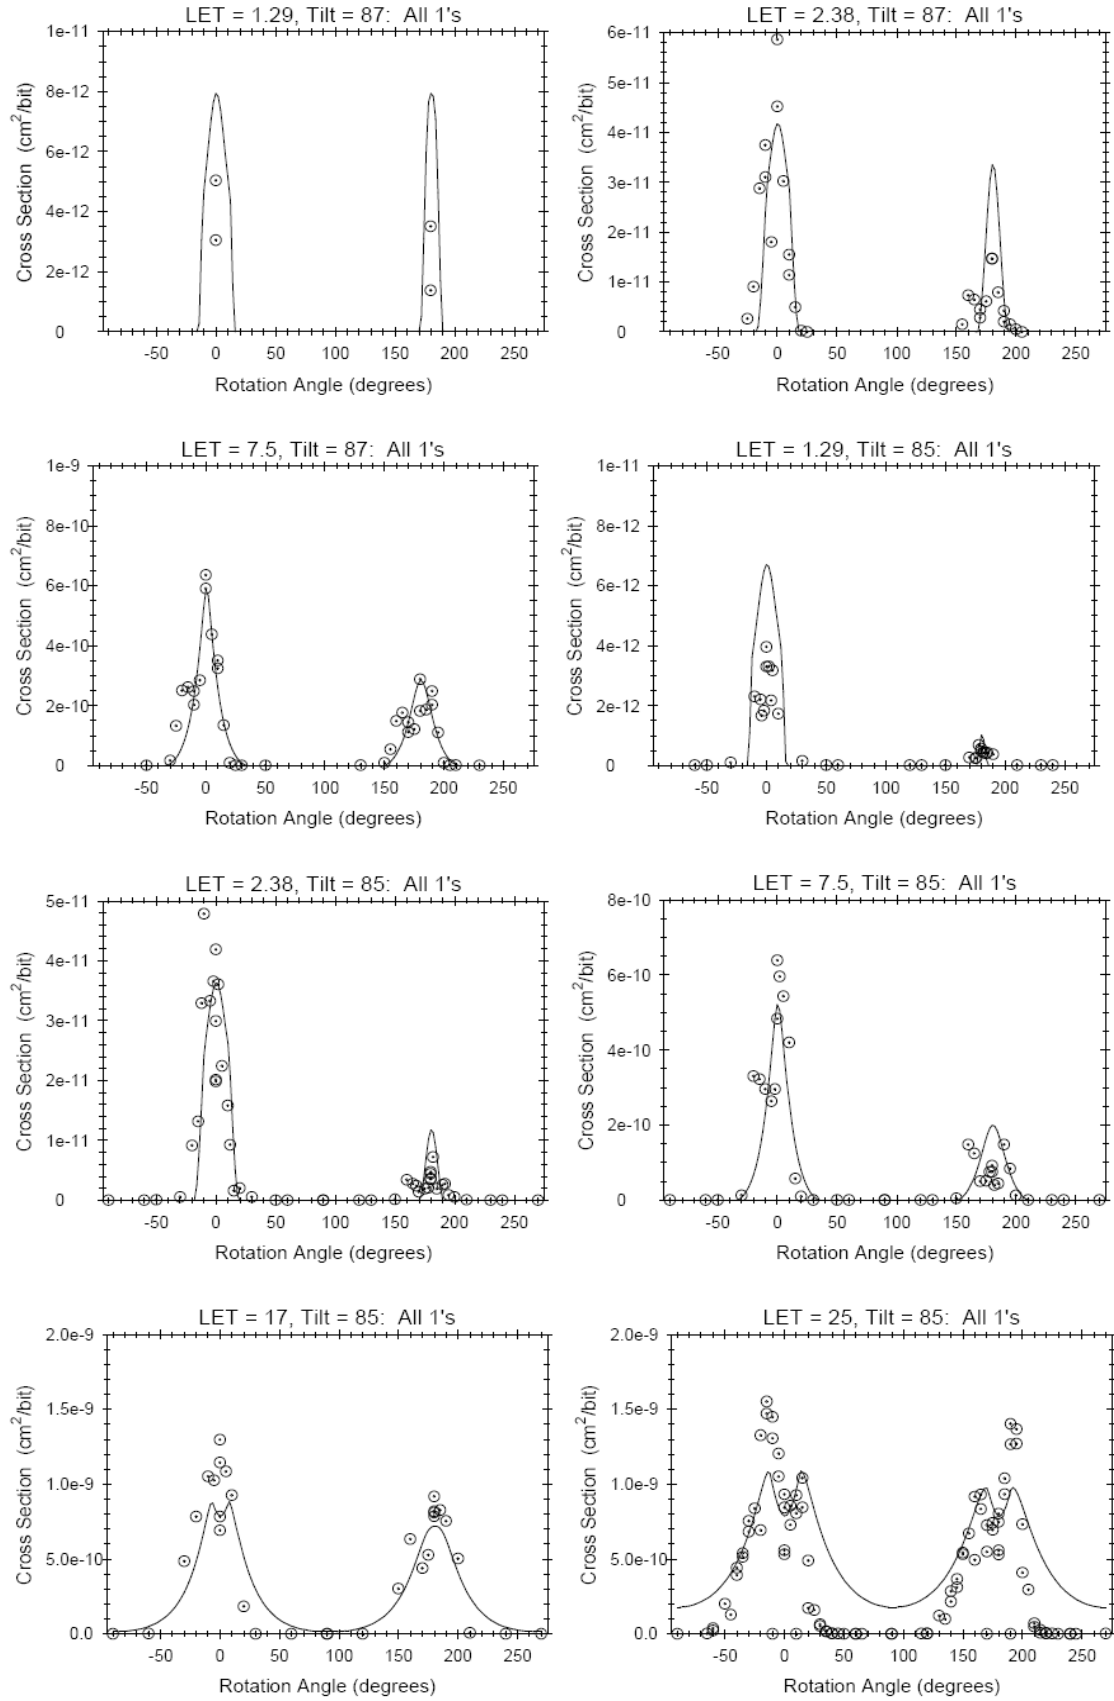
The RHBD configuration cell used in the Virtex-5QV has internal redundancy so that, if any single node collects charge, it will not upset, although it is possible that there may be a brief transient on the cell output. Indeed, unless two nodes collect at least the minimum charges  $Q_{crit1}$  and  $Q_{crit2}$ , the cell will not upset. The pairs of nodes that can induce upset by simultaneously collecting charge are intentionally spaced a good distance apart. This results in an upset susceptibility for a given ion that varies widely (a few orders of magnitude) depending on the ion's direction vector. The most sensitive direction has the ion vector aligned with the straight line between the two nodes.

### 4.1 Test Chip Static Results

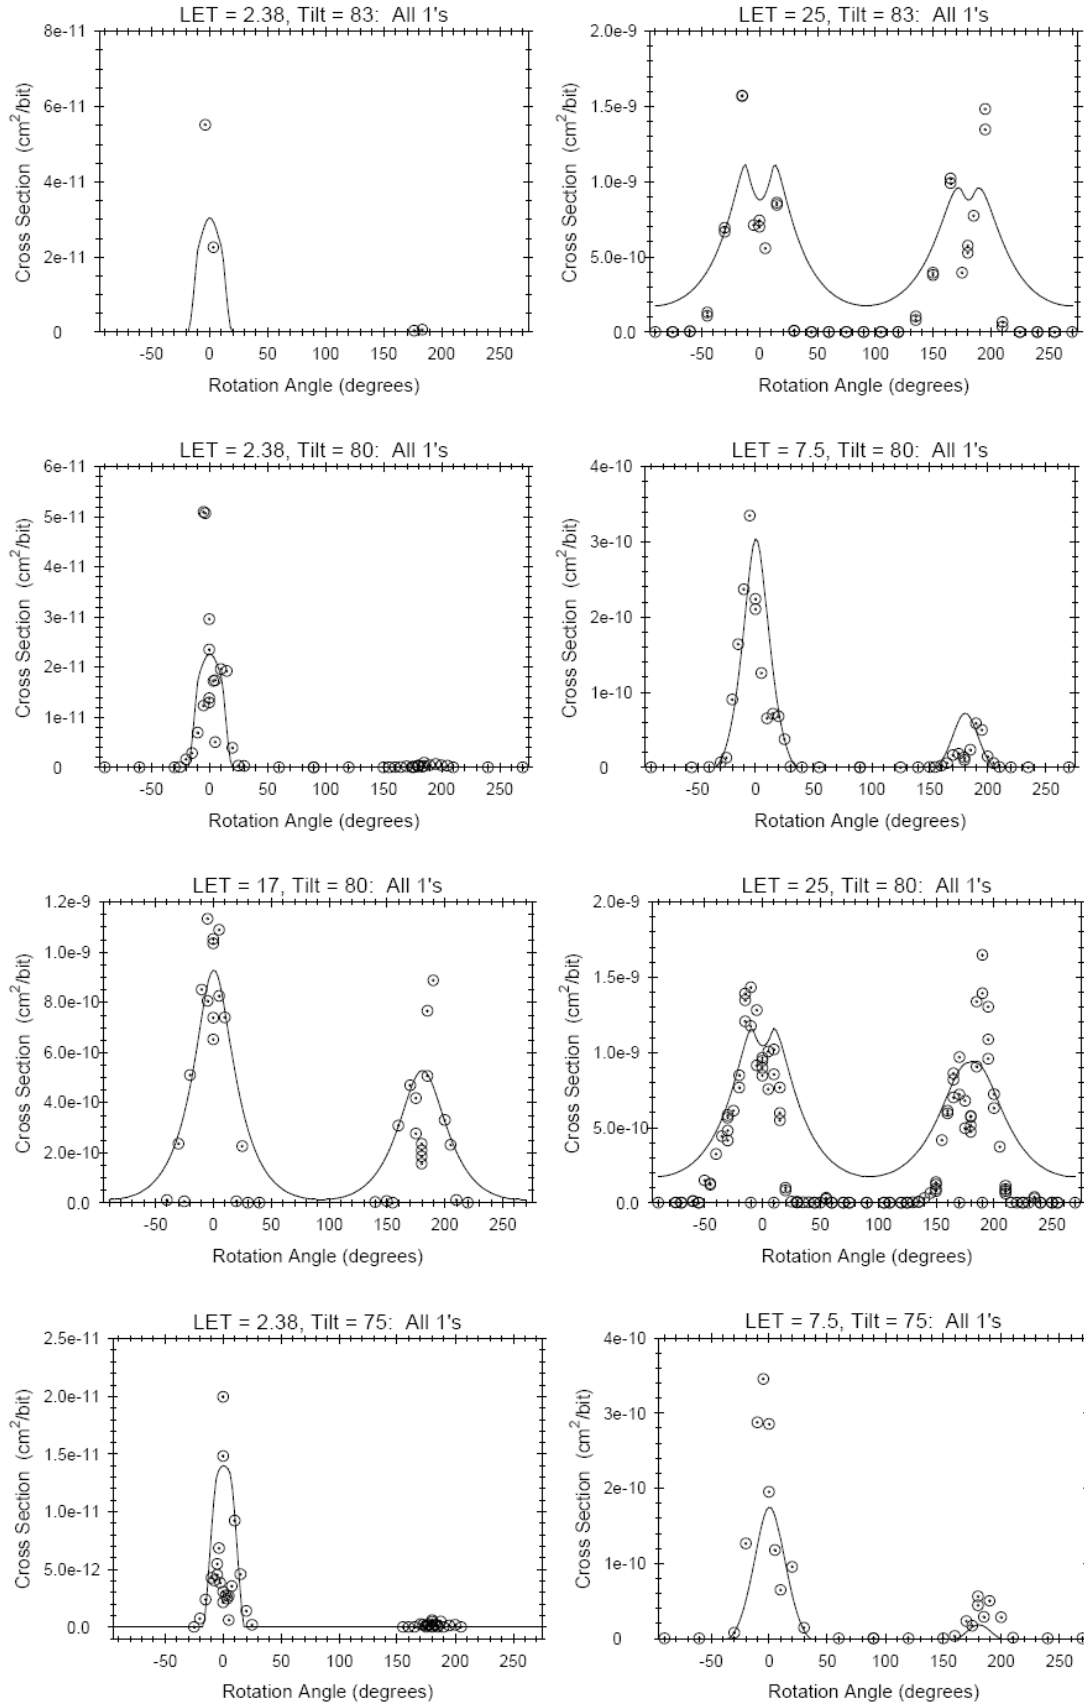
Because the most sensitive direction goes through the pair of active nodes, the most sensitive angles of incidence are in the plane of the “top” of the silicon, i.e., grazing angles. Experimentally, this is a problem as the apparent thickness of any intervening dead layer increases rapidly (with the cosine of the angle) near those grazing angles. Available accelerator beams have limited penetration depth and, if the angle is too steep, will not reach the active silicon layer. Thus, two painful conclusions arise:

1. It is impossible to measure at the most important angles, and
2. Using a flip-chip device (even with aggressive backside thinning) exacerbates this problem.

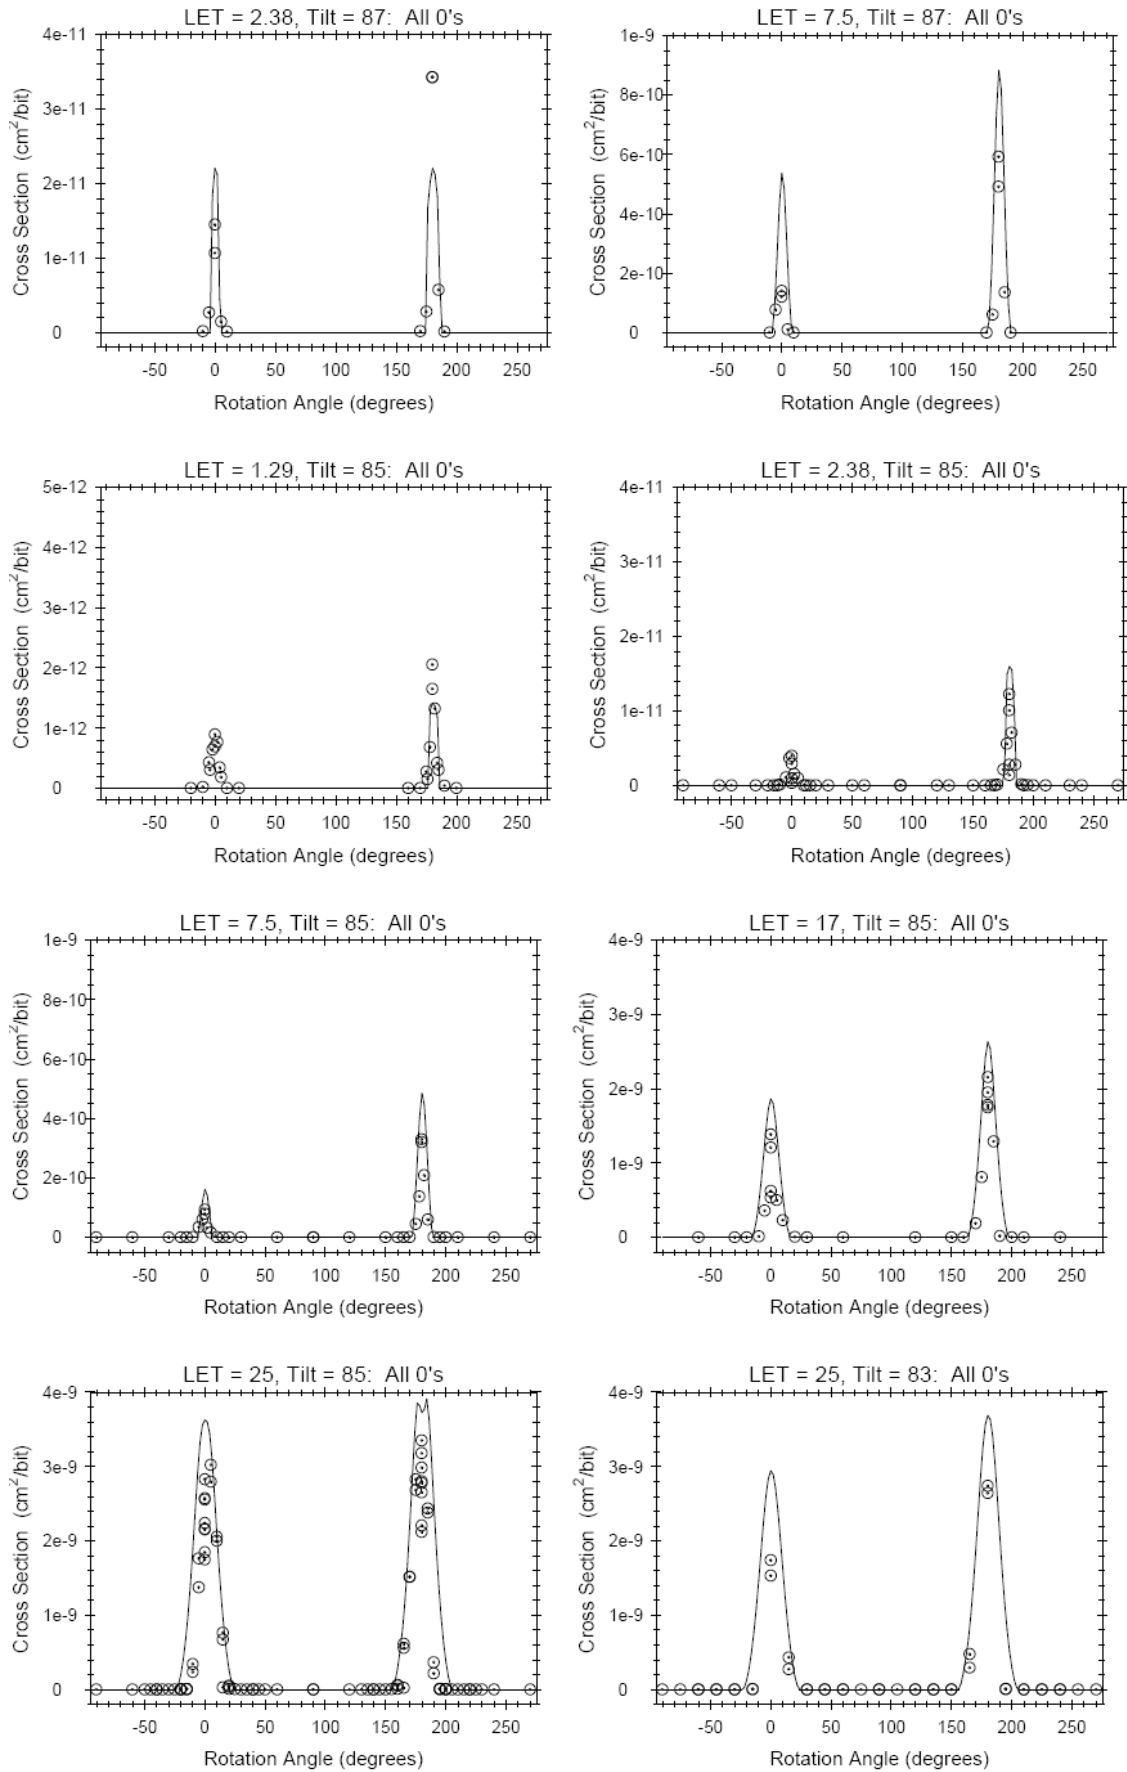
Still data from a wire-bond test chip can get to steeper angles than a thinned flip-chip FPGA. Therefore, an experimental data set has been taken using a face-up test chip incorporating an array of the configuration cells. Sample plots of that extensive dataset are given on the following pages. The circles are the measured data points and the solid lines are the fit to the Edmonds dual-node physical model described briefly below in Sec. 4.4 and more extensively in Ref. 12. Note the test chip structure for the configuration cell data set shown in Figures 13 and 14 is not exactly the same as the final layout, but is shown because it is the most complete dataset and, thus, was used to explore the model's fit and sensitivity. The cell's layout was modified very slightly at the last minute and a distinctly smaller dataset has been taken to date on test chips with the actual product cell. Edmonds notes that the difference in SEU response between the actual production configuration cell and the cell with the extensive dataset appears to be small [12].



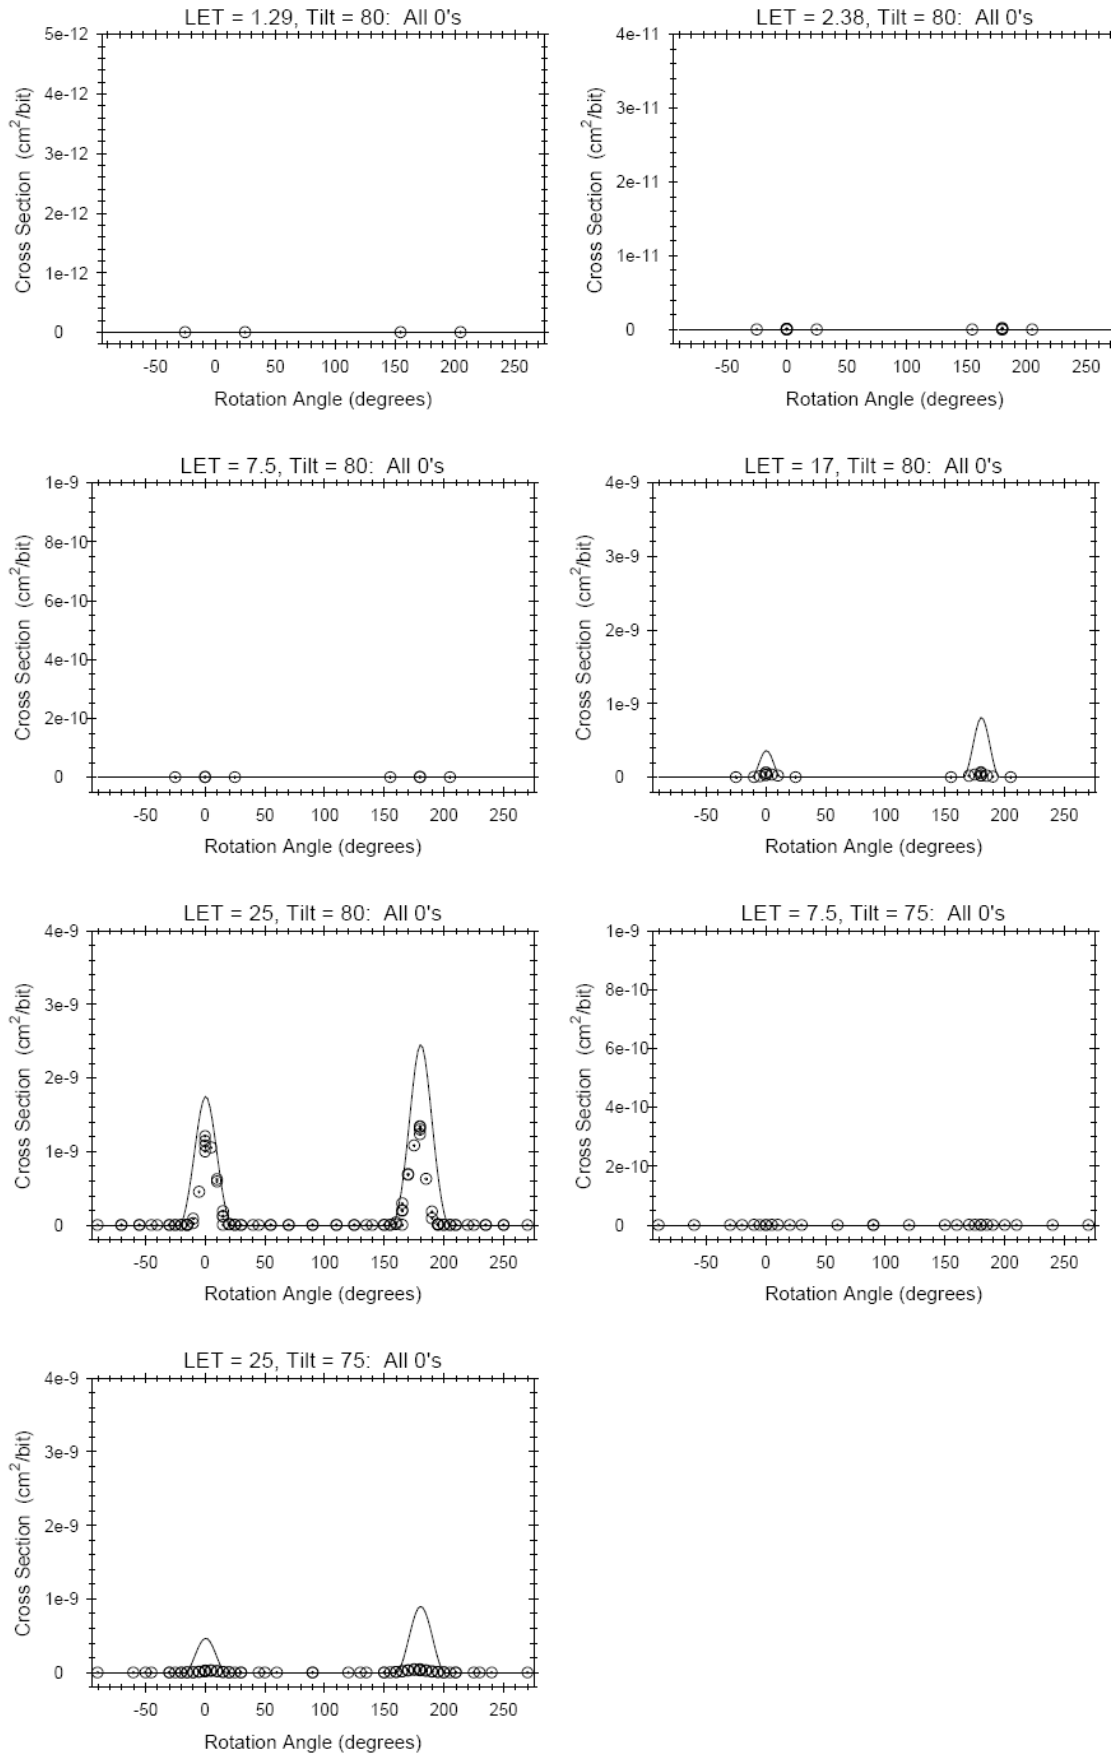
**Fig. 13a.** Rotation sweeps for dual-node configuration cells storing ones [Ref. 12, Fig. 23 pt 1]



**Fig. 13b.** Rotation sweeps for dual-node configuration cells storing ones [Ref. 12, Fig. 23 pt 2]



**Fig. 14a.** Rotation sweeps for dual-node configuration cells storing zeros [Ref.12, Fig.28 pt.1]



**Fig. 14b.** Rotation sweeps for dual-node configuration cells storing zeros [Ref.12, Fig.28 pt.2]

## 4.2 Configuration Cell Direct Upsets from Protons

In an investigation led by Munir Shoga of SAIC, and Gary Swift, the lower LETs of the steep angle dataset help rule out any significant contribution to direct ionization upsets from protons. For example, 40 MeV/amu N (with an incident LET calculated to be  $\approx 0.72$  MeV-cm<sup>2</sup>/mg) at 85 degrees in alignment with the dual nodes yielded a statistically significant number of upsets and a (non-effective) cross section of about  $2 \times 10^{-13}$  cm<sup>2</sup> per bit. Although a proton can directly deposit more than an LET of 0.7 as it slows down near the Bragg peak, it cannot do so to two nodes with any distance between them. Also note that the cross section is about as low as that obtained from proton-induced reactions in single-node devices; dual-node devices get an additional “geometrical” factor from having a narrow acceptance cone.



Photo 2. Vacuum chamber setup on the beamline at UC-Davis cyclotron (left) and the Virtex-5QV DUT at 75 degrees with dual nodes aligned with beam direction (right).

Earlier (in June 2010), a direct ionization upset test was run in vacuum with protons (see Photos 2a and 2b) at low energies at 75 degrees and with the rotation angle such that the dual nodes were aligned with the beam. The low energies were carefully selected to approach and bracket the region where the Bragg peak in energy deposition was in the device’s thin epitaxial layer. At 980 keV incident energy, five BRAM upsets were observed indicating a cross section of about  $2.3 \times 10^{-19}$  cm<sup>2</sup> per bit and, at 3.2 MeV, the observed BRAM cross section was almost 5 orders of magnitude bigger at  $1.1 \times 10^{-14}$  cm<sup>2</sup> per bit. At the lower energy, the protons are likely stopping short of the active region with a few stragglers getting through. At the higher energy, direct ionization upsets of the BRAM cells seem clear with a cross section approaching that seen from indirect upsets at high energy. No configuration cells were upset during this test.

In conclusion, these results indicate that the 65nm dual-node configuration cell used in the Virtex-5QV is fully immune to direct ionization upset from protons.

### 4.3 SET-triggered FPGA Results

Using upset-hardened-by-design techniques, it is possible to drive the direct upset rate down to such a low level that SET-triggered upsets dominate. In the static case, SETs coincident with a clock edge are ruled out, but SETs on asynchronous control lines (say reset or write signals) or on the clock lines themselves can cause an upset.

Two main categories of SET-triggered upsets of the dual-node configuration cells have been observed:

1. triggered events in “capture” cells, and
2. unintended “writes” to configuration cells.

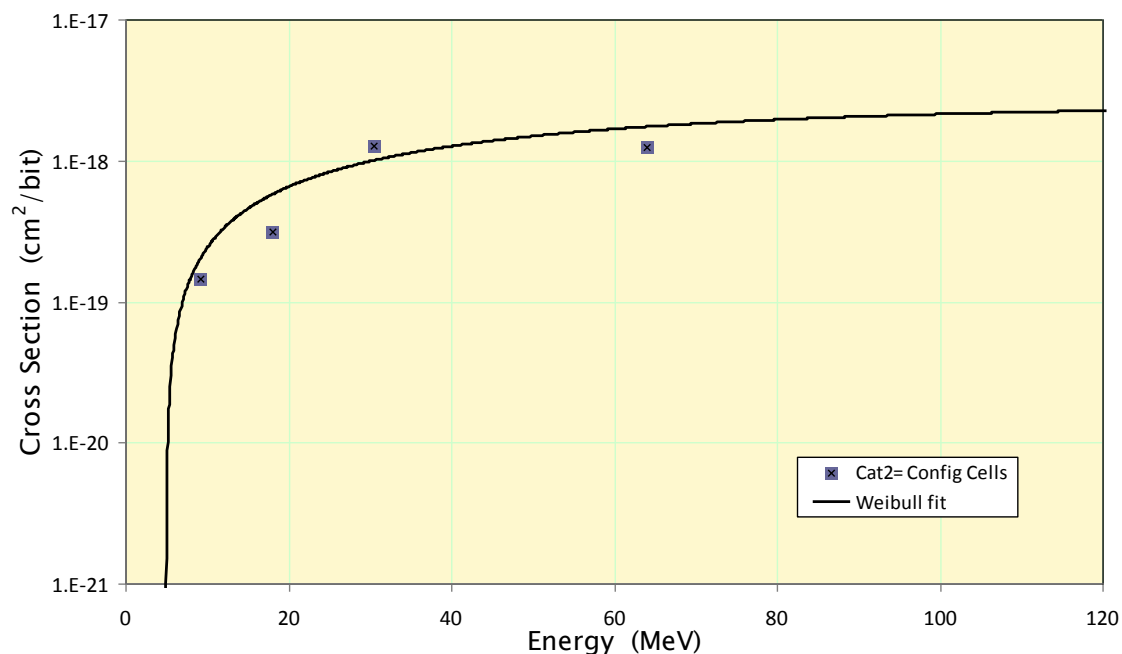
Category 1 cells are not truly configuration cells in that they don’t control or route anything; as a result, whether they are upset or not is irrelevant to correct design operation. However, upsets will add to the “false alarm” rate for detection schemes because Category 1 bits do appear in the readback bitstream.

On the other hand, category 2 upsets may break a design and so definitely add to the static cross section of the FPGA a component that adds to what was seen on the test chip.

Real-time sorting and masking capabilities of the Configuration Monitor allows the XRTC test apparatus to collect data on both these SET-induced upset types simultaneously with any other testing of the Virtex-5QV. More than 100 hours of testing in the July and October 2011 TAM campaigns yielded heavy ion data for defining these cross section vs. LET curves, but the analysis is still proceeding. Similarly, proton data has been collected in the November 2011 and January and June 2012 UC-Davis tests. The expectation is that the responses will be the standard single-node type and, thus, can be fit with Weibull curves that can be fed into CRÈME. The analysis is a bit painful, but the result for the category 2 upsets for LET=88.7 (normal incident gold ions) is encouraging. The event cross sections are  $4 \times 10^{-11}$  and  $7 \times 10^{-12}$  cm<sup>2</sup> for ones and zeros, respectively. Thus, it is likely that they will be overshadowed by the direct upset rate which is less than 5 per year in GEO (see the next section).

A detailed manual analysis of effective LET=145 (gold at 55 degrees) shows event cross section for category 2 climbs about a factor of six to  $2.4 \times 10^{-10}$  cm<sup>2</sup>/bit. For the whole device, the category 1 (capture bits) event cross section is pretty close at  $8.3 \times 10^{-3}$  cm<sup>2</sup>/device. On a per bit basis, this is a much bigger cross section as there are fewer than 200,000 capture bits, but over 30,000,000 ‘real’ configuration bits. Interestingly, for both categories (and in the absence of the huge capture bit ‘clobbers’ of thousands of bits), the cells storing zero rarely show events with more than one upset while cells storing one have an average event size of three to five. For category 2 upsets (inadvertent writes), this is likely the result of the predominance of zeros; zero is much more likely to be the value on the bus when an SET induced write event happens and, in that case, stored ones are much more likely to be upset. Under this explanation, writing a zero into a cell storing zero is, by far, the most likely event happening, but, of course, that event looks exactly like no event. Note that another implication of this explanation is that there is a design dependent component to the experimental results: the ratio of ones to zeros in the configuration matters. For this data point, the “heater” design used for latchup testing was the design-under-test.





**Figure 15.** Virtex-5QV Configuration Bit Upset Susceptibility due to Protons

Preliminary proton results and Weibull fit for category 2 errors are shown in Figure 15. These results are extracted from the data taken at UC-Davis in November 2011 and January and June 2012 and are ‘raw’ results as upset rates and event rates are not yet separated. This is probably not too important as it appears the event rate will not be substantially lower than the upset rate except perhaps for cells storing ones. The device cross sections for category 1 upsets is more than an order of magnitude larger even though the number of bits is over two orders of magnitude smaller. To reiterate, upsets of category 1 cells do not affect design operation; however, they will give false alarms on configuration upset detection unless they are filtered out.

#### 4.4 Orbital Rate Calculations

There is no industry-accepted model of how to turn a set of dual-node susceptibility measurements over angles into rate predictions for particular space radiation environments. Larry Edmonds of JPL has considered this problem for a long time and has developed a physical model (albeit with ten fitting parameters) to overcome the necessary extrapolation problem. He has written a formal JPL report [12] that contains the model's derivation and its application to the Virtex-5QV test chip data set. He concludes that the solar minimum GEO rate behind 100 mils of Al-equivalent shielding is  $2.1 \times 10^{-10}$  and  $5.3 \times 10^{-11}$  upsets per bit-day for ones and zeros, respectively.

Interestingly, he also obtains equivalent single-node Weibull parameters that give approximately the same answer for the same environment and he does recommend using these to estimate the rates in other space environments. This is basically averaging a very peaked rotation angular distribution over all rotation angles so that an equivalent single-node device is used for projecting rates. The Weibull parameters for the equivalent device are copied from Reference 12 and given in Table 9. For GEO, they give  $1.6 \times 10^{-10}$  and  $4.2 \times 10^{-11}$  upsets per bit-day for ones and zeros, respectively, or about 25% lower than the answers from the full model.

**Table 9.** Virtex-5QV Heavy-Ion Weibull Fit Parameters for an Equivalently Hard, Single-Node Configuration Cell

Cells	Weibull Parameters			
	Limit, cm <sup>2</sup> /bit	Onset, MeV-cm <sup>2</sup> /mg	Width -	Power -
storing "ones"	$3.1 \times 10^{-9}$	0.20	61	2.55
storing "zeros"	$8.0 \times 10^{-9}$	0.25	103	2.95

## 5 CONCLUSION

The Virtex-5QV, a megarad(Si) RHBD SRAM-based reconfigurable FPGA designed for space, performed well in these heavy-ion and proton irradiations, exhibiting no single-event latchup (SEL) even at elevated temperature and spec-max voltages to effective LETs above 130 MeV cm<sup>2</sup>/mg and fluence above 10<sup>8</sup> per cm<sup>2</sup>.. The XQR5VFX130 exhibited extremely low total SEFI susceptibility to heavy ions and protons with a resulting rate of approximately one per 10,000 years in the geosynchronous orbit's (GEO) radiation environment. Upsets of the unhardened Block RAMs might be a significant concern based on these data; however, the enhanced ECC capabilities built into the FPGA are more than sufficient to maintain error-free design operation; in-beam testing has proven their effectiveness [8]. The projected configuration upset rates (a few bits per year in GEO) are likely acceptable for most applications; based on a fairly conservative architectural vulnerability factor (AVF) of 10x, an actual error (circuit malfunction) is expected to occur less than about once every two years of on-orbit operation on average. The most critical applications can use selective TMR to operate error-free in spite of the presence of an upset. Employing design-level mitigation reduces the system error rate due to upsets to well below that of the SEFI rate for even the worst-case space environments.

## 6 REFERENCES

- [1] *Virtex-5QV FPGA Packaging and Pinout Specification*, UG520,  
[http://www.xilinx.com/support/documentation/user\\_guides/ug520\\_V5QV\\_Packaging\\_and\\_Pinout.pdf](http://www.xilinx.com/support/documentation/user_guides/ug520_V5QV_Packaging_and_Pinout.pdf)
- [2] Overview Datasheet,  
[http://www.xilinx.com/support/documentation/data\\_sheets/ds192\\_V5QV\\_Device\\_Overview.pdf](http://www.xilinx.com/support/documentation/data_sheets/ds192_V5QV_Device_Overview.pdf)
- [3] *Radiation-Hardened, Space-Grade Virtex-5QV FPGA Data Sheet: DC and Switching Characteristics*, DS692,  
[http://www.xilinx.com/support/documentation/data\\_sheets/ds692\\_V5QV\\_Data\\_Sheet.pdf](http://www.xilinx.com/support/documentation/data_sheets/ds692_V5QV_Data_Sheet.pdf)
- [4] G. Allen, G. Swift, C. Carmichael and the XRTC, *Virtex-4QV Static SEU Characterization Summary*, Sept. 2008,  
[http://parts.jpl.nasa.gov/wp-content/uploads/NEPP07FPGA\\_v4Static.pdf](http://parts.jpl.nasa.gov/wp-content/uploads/NEPP07FPGA_v4Static.pdf)
- [5] G. Swift, ed., *Xilinx Single Event Effects Consortium Report: Virtex-II Static SEU Characterization*, Jan. 2004,  
[http://parts.jpl.nasa.gov/wp-content/uploads/virtex2\\_010411.pdf](http://parts.jpl.nasa.gov/wp-content/uploads/virtex2_010411.pdf)
- [6] *Virtex-5QV FPGA Quick Start Guide*, XMP120.
- [7] Y.C. Wang, *Virtex-5QV FPGA External Configuration Management*, XAPP588,  
[http://www.xilinx.com/support/documentation/application\\_notes/xapp588\\_v5qv\\_ex\\_c\\_onfig\\_mgmt.pdf](http://www.xilinx.com/support/documentation/application_notes/xapp588_v5qv_ex_c_onfig_mgmt.pdf)
- [8] G. Allen, L. Edmonds, C.W. Tseng, G. Swift, and C. Carmichael., "Single-Event Upset (SEU) Results of Embedded Error Detect and Correct Enabled Block Random Access Memory (Block RAM) within the Xilinx XQR5VFX130," *2010 Radiation Effects Data Workshop Record*, <http://parts.jpl.nasa.gov/wp-content/uploads/EDAC-final-IEEE-proof.pdf>

- [9] L.W. Connell, P.J. McDaniel, A.K. Prinia and F. W. Sexton, “Modeling the heavy ion upset cross section,” *IEEE Trans. Nucl. Sci.*, vol 42, No. 2, April 1995.
- [10] R. Monreal, *Single Event Effects Test Report: Virtex-5QV Field Programmable Gate Array-Digital Signal Processors*, SwRI Document No. 14520-RATR-03, July 2011, <http://parts.jpl.nasa.gov/wp-content/uploads/14520-RATR-03.pdf>
- [11] CREME website: <https://creme.isde.vanderbilt.edu/>
- [12] L. D. Edmonds, *Estimates of SEU Rates from Heavy Ions in Devices Exhibiting Dual-Node Susceptibility*, JPL Publication 11-6, June 2011, <http://parts.jpl.nasa.gov/wp-content/uploads/Dual-Node-JPL-pub1.pdf>

## 7 APPENDIX

The single-node, heavy-ion SEU response data sets have been fit with Weibull curves to facilitate orbital rate calculations. The fitting equation is:

$$\sigma(LET) = \sigma_{sat} (1 - \exp \{ - [(LET - L_{th}) / W]^S \})$$

$\sigma_{sat}$  is the limiting or plateau cross section (or “limit”),

$L_{th}$  is the LET threshold parameter (or so called “onset”),

W is the width parameter, and

S is a dimensionless exponent dubbed “power.”